GEMFIRE®
the enterprise
data fabric

High Performance, Mission Critical Applications for the war-fighter:

Solutions to network challenges and today's fluid combat environment

There are key challenges in building mission critical applications that must work in the face of the following conditions:

- **Data access when data sources (especially multiple) are not collocated with or near the application (increasing latency)**
- **Times/situations when the network or data sources or both are unavailable**
- **Low bandwidth and/or unreliable networks**
- **Data that is changing at high rates and these changes need to be made visible to multiple applications rapidly**
- **Automatic data synchronization between applications and data sources when the application is cut off from the network and then returns to the network**
- **Hardware, software, network failures**
- **Combine data from multiple data sources and create efficient data placement**
  - **o move data by interest**
  - **o rollup data from multiple sources**
  - **o drill down to details on as needed**

To enable agile, robust, dependable information, it must be available when it is needed at speeds that enable it to matter. With Net-Centric systems, the network is the database. To avoid having the network become the problem, data must be delivered so as to meet these goals. Recently accessed data should be available even if the network is off-line. Data should be as current as the network allows.

Information is a force multiplier. Accurate and timely information is a bigger force multiplier. With the advent of GIG-BE, war-fighters have come to expect fast and reliable access to information in garrison and command centers. Unfortunately, network access erodes as war-fighters move to more mobile environments, ships, mobile command centers, special-forces with sat radio, etc. Mobile environments have significantly degraded capabilities to deliver information affecting timely information access. It can be frustrating and lethal not to be able to access information that was available only minutes before when a network connection vanishes. The ability to subscribe to and receive continuous updates at a fine grained level to key information can enable planners and operators to get the details that can cause them to change their minds as the situation changes. With today's much more fluid combat environments, high level commanders need access to more tactical details than ever before. And they need those details immediately.

If it was only possible to hang onto information that was being looked at minutes before, or to browse information in garrison or on the network and have it travel with the war-fighter as they deployed. If only it was possible for that information to be automatically updated whenever that person was on the network. If only it was possible to subscribe at a fine grained level to key information and get alerted as data that affected an operation changed, operators could be aware of the details that planners depended on changed. And if only it was possible to take tactical real-time feeds, to aggregate them and roll them up to high levels, yet have ability to drill down to the detailed information if necessary.

The simplest use case is to maintain access to critical data in the face of network outages. For applications such as the NCES Mediation Service or the IA Security service applications, simple data caching can enable access to data that they have been working on even though the network connection is down. Services would only look in the data cache if they were unable to reach the source data over the network. This could be because the network is down or the service is down or the data source is down.

The next level of capability is to reduce the latency and network bandwidth usage. Simple caching is insufficient to solve this issue for data that has high rates of change or for which changes have serious impact to the application, because the cache may be out of sync with the actual data source. However, if a service or data source actively send changes to the cache(s) (real-time updates when network is available and guaranteed sequential delivery of data change events after a network outage) then it is possible to:

- Get data in face of network outages
- Reduce network band usage
- Speed data access

If the cache supports notification of data changes back to the application, then it is also possible for applications to be aware of data changes to the data they have previously retrieved (i.e. do automatic view maintenance). Most applications today are not written to expect this level of service. But it can have a significant benefit to be aware that data in the application is stale.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
APR 2010		**00-00-2010 to 00-00-2010**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
High Performance, Mission Critical Applications for the war-fighter: Solutions to network challenges and today's fluid combat environment	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **GemStone,1260 NW Waterhouse Ave., Suite 200,Beaverton,OR,97006**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
Presented at the 22nd Systems and Software Technology Conference (SSTC), 26-29 April 2010, Salt Lake City, UT. Sponsored in part by the USAF. U.S. Government or Federal Rights License

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**43**	

This still doesn't solve the issues of low bandwidth network and high data change rates. So the next level of capability would be a caching technology that supports those capabilities necessary to only send data of interest to applications, so that unnecessary bandwidth is not spent sending data that is not yet of import to the user. For example, the GCCS-J TMS application is able to manage 100,000 tracks with a 1 minute update rate (which works out to approximately 1,650 update per second of 1k objects or 1 MB per second). With thousands of concurrent users that totals Gigabytes per second. Most users can only look at thousands of object of interest (and zoom down to even less when analyzing a situation such as air defense cover of the straits of Tehran). So by being able to filter the data and detect the delta, bandwidth usage can be reduced an order of magnitude or more. Additionally, if many application users sit at a location, then the ability to tier software so that local server caches can be shared by many local user caches can also significantly reduce bandwidth usage. In this case it can also reduce latency for the 2nd, 3rd and nth local user of that data. For example, in the case of a carrier at sea, sending the same video clip, or thumbnail image to multiple users, results in that data traveling over a very finite pipe more times than is necessary. Ideally, a solution that handled these issues would allow for multiple expressions of fine grained interest, so that if users were interested in multiple subsets of data that had different characteristics they could get them all. e.g. Enemy air tracks and friendly air defense units. Any solution to handle these issues must be capable of scaling to handle any arbitrary rate of data change (so long as the necessary network to send the data is available).

Interest in data should be able to last for as long as necessary and should not actually require a human being sitting at a screen. Imagine a mission planning system that retrieves data from a number of sources such as tracks, order of battle, and battle damage assessment. In the course of planning the mission, new information could arise, a target gets destroyed, a new enemy SAM unit shows up, which would cause different decisions during planning. Even after the mission started, knowledge of changes to the underlying data could be distributed to participants that needed it. Once data from multiple sources with fine grained detail can be pulled into an application, new questions can be asked of the data not anticipated by the designer of either source. For example, merge MIDB with 'contains' information about equipment with track data and ask for all current tracks of EW capable aircraft.

Next are the issues of synchronization when applications that were cut off from the network and are reconnected at a later time. Data synchronization falls into two categories, easy (in that it does not have to know about application context), and hard (application context matters because of potential conflicting updates when the application was off-line). Although the "easy" data synchronization problems involve lots of complex technology to automate the synchronization process, they do not require any work from application developers and are thus deemed easy. The "hard" data problems require only a little technology, but require a lot of domain knowledge and work from application developers, and are thus deemed hard. (At least for the people actually building an application).

Where multiple users can not update the same data in a distributed environment, the servers need to queue data changes to individual clients based on their interests. This may seem a simple challenge but for fast changing data this can create a queue of significant size that exceeds the client's and network's ability to ever have the client catch up. In that case a mechanism that conflates changes so that the client only sees the most recent update to each piece of data would minimize the overall size of the queue and maximize the chance of success. If message ordering is to be preserved for all clients, then queuing of messages is also necessary for client applications that fall behind in their ability to process incoming data because they are too busy. If a client "never" recovers for any of the above reasons, then it may be desirable for performance reasons to be able to automatically remove their queues and data interests until they reemerge. By "never", we really mean that a client falls more than some configurable amount of time or number of messages behind the overall system.

Where multiple users can update the same data, the caching mechanism must have a pluggable way of resolving conflicts. Since a simple last one to update the data "wins" does not work for most applications, application specific code that checks originators, timestamps, version numbers, or other application specific data, must be able to be applied to incoming data to apply correct changes and deny incorrect changes to data. Notification of failed updates because of data conflicts must be propagated back to the original application instance that tried to update the data. A mechanism that validates incoming data can also be used to merge, roll up, or otherwise combine multiple data instances in the data fabric into a single element. This too would depend on application specific code that is invoked as data arrives.

Because this paper is focused on describing potential solutions to challenges for mission critical applications, it is important to consider failure conditions. So far all of the potential failures described have been network issues or edge client issues. Preventing data loss and assuring continuity of operations in the event of server hardware loss is also critical. To assure data integrity across multiple instances of the same application, messages to clients must be delivered even in the event of a server loss. To keep up with data rates for high speed data, this often must be done without committing data to disk. It is critical to replicate 'in memory' the server caches and queues of messages

waiting to be sent to clients that have fallen behind or are off the network. This data should be able to be replicated on N nodes, (where N is the number needed for the application to reach the desired level of redundancy to assure continuous operation). Having done this, it is possible to have client applications fail-over seamlessly from one server to another in the event of server loss. Such a mechanism can even accommodate site failure, although depending on data rates, clients may need to resynchronize with the servers. The ability to fail-back as servers become available is also critical to maintain load balanced and guarantee SLAs.
If distributed caching is used to manage the current operational data. Some high bandwidth locations, the Command HQ's, the Pentagon, etc, could combine data from multiple servers and roll that data up into a larger operational picture. Only the higher level rolled-up data need be sent up the command chain. Since the details still exist at each of the individual servers where the data came from, and since the data can be queried in such a way that deltas are continuously streamed back to an application, then the larger, more coarse grained operational picture has the ability to enable commanders to drill down for live details when they need them. This enables a virtual total picture to be available without saturating the available bandwidth.

In order to have the ability to solve the challenges of high data rate mission critical applications that must work in the face of various network issues, be able to filter data, allow for occasionally connected applications, handle synchronization, be tolerant of hardware failures, and scale up or down as needed, a new class of software infrastructure is needed. This software must be able to provide for all of the capabilities described in this paper including dynamic data filtering, high data rates, handling slow receivers, handling disconnected users, managing network or hardware loss, preventing data loss, have mechanisms to ensure data integrity and the ability to scale to data rates as needed, in massively distributed environments. Such a type of software, a "Data Fabric" is available today from GemStone Systems. It is being used by 7 of the top ten Wall Street investment banks for program and algorithmic trading, risk and data analysis, global order book management (across WAN connections), global market and reference data distribution of original and derived data to applications and user around the globe. It has been put into a next generation GCCS-J/COP prototype.

GemFire is a high performance, distributed data management infrastructure. It provides distributed in-memory data caching, application level notification of data changes, configurable ACID properties, query, and enterprise scale. It supports multiple topologies including tiered servers, with edge clients. It supports fault tolerance for hardware and network failures at either the server or client level. This enables support for disconnected operations. Its memory management abilities enable it to be configured to replicate data for high concurrent loads, and for high availability. Data in memory can be grouped into separate Regions to support data placement and to support separating data from multiple applications. The distributed memory management also provides for methods to gracefully expand capacity to meet scalability and performance goals. GemFire has been benchmarked at over 100,000 1k inserts per second while distributing data to 1,000 client VMs. In general it is limited only by network bandwidth and available hardware. In order to provide for the ability to maximize performance and also support the level of concurrency needed for individual distributed applications, GemFire supports both high coherence and loose coherence across the distributed memory space. This included support for distributed transactions in a peer-to-peer environment.

The ability of applications to register interest in data with fine-grained queries, using OQL, a superset of SQL 92, from the Object Data Management Group (ODMG), enables applications and users to dynamically have the data they need move to clients from one or more servers. Applications receive a result set from the query. Then applications continue to receive notification as data changes on the server change the query results. Updates, deletes, and inserts are sent to the client result set to provide automatic view maintenance. Client applications receive a call-back and can take appropriate actions in response to the data changes. This ability like all of the data delivery mechanisms in GemFire is fault tolerant. Servers have the ability to automatically queue data for clients that are slow receiving data or are currently off the network. Data and queues can be replicated across multiple servers. If a server dies, clients seamlessly switch to an alternate server and continue receiving data from where they left off.

GemFire also provides for the ability to run regular queries that do not continue to send data changes to the application. Regular queries can be written to bring back less than the full object. For large objects, such as Air Tasking Orders, which can exceed 40 MB, this ability can significantly increase performance and reduce network bandwidth, when users are only interested in a small subset of the data.

GemFire provides for data persistence. This combined with the continuous query capabilities enable the easy creation of client applications that can be shutdown, moved, restarted, and periodically brought on the network to send or receive data changes. For the issues of conflict resolution that emerges when data can be updated at multiple locations GemFire provides multiple mechanisms to plug in application defined code to handle conflict resolution. There are two major ways to handle this issue. The first is to have updates travel via a separate GemFire Region to a process that is used to scrub data inserts. For the GCCS-J TMS-E application prototype, this method is used as all user changes, pass back into the correlation engine. The other methods is to have server-side plug-ins, called cache-writers,

compare the input data to the current state of the data, and then decide which data state is correct, using time-stamps, authoritative source, version numbers, or other application specific data. If this method needs additional reference data to make the correct decision, this data could easily be stored inside GemFire in another Region. Since client application block when inserting data into GemFire, changes that are rejected are immediately visible to end-users. Note, this method only works when users are on-line. For off-line changes, data must flow to an alternate region where data can be examined using the first method.

GemFire also provides the enterprise scale management capabilities needed deploy large groups of distributed caches. It is highly instrumented. Its capabilities include distributed management, monitoring, alerting, and administration of the distributed system. All of these capabilities are exposed through the industry standard management API's JMX. It also provides for retrospective analysis of performance. It is highly configurable and tunable over with over 100 configuration parameters.

Given NCES's vision statement to "enable the secure, agile, robust, dependable, interoperable data-sharing environment for DOD where war fighter, business, and intelligence users share knowledge on a global network", a distributed caching capability that provides for disconnected operations, high speed data distribution, notification of data changes, and enterprise level fault tolerance and management could provide significant capabilities to further this vision.

Corporate Headquarters:
1260 NW Waterhouse Ave., Suite 200 Beaverton, OR 97006 | Phone: 503.533.3000 | Fax: 503.629.8556 | info@gemstone.com | www.gemstone.com

Regional Sales Offices:
New York | 5 Penn Plaza, 23rd Floor New York, NY 10001 | Phone: 646.530.8458
Washington D.C. | Phone: 301.564.0550

GEMSTONE®

Copyright© 2008 by GemStone Systems, Inc. All rights reserved. GemStone®, GemFire™, and the GemStone logo are trademarks or registered trademarks of GemStone Systems, Inc. Information in this document is subject to change without notice.

GEMSTONE®

Enterprise Data Fabrics (EDF)

Michael Lazar, Federal Technology Director
Mike Sinisgalli, Director Public Sector

THE GLOBAL #GRIDAWARDS 04-28-08
Best Vendor Grid Solution

GEMFIRE®
the enterprise data fabric

Agenda

GEMSTONE®

- Key Trends
- The Challenge
- Enterprise Data Fabrics
- Use Cases

GEMFIRE®
the enterprise data fabric

Exponential growth

GEMSTONE®

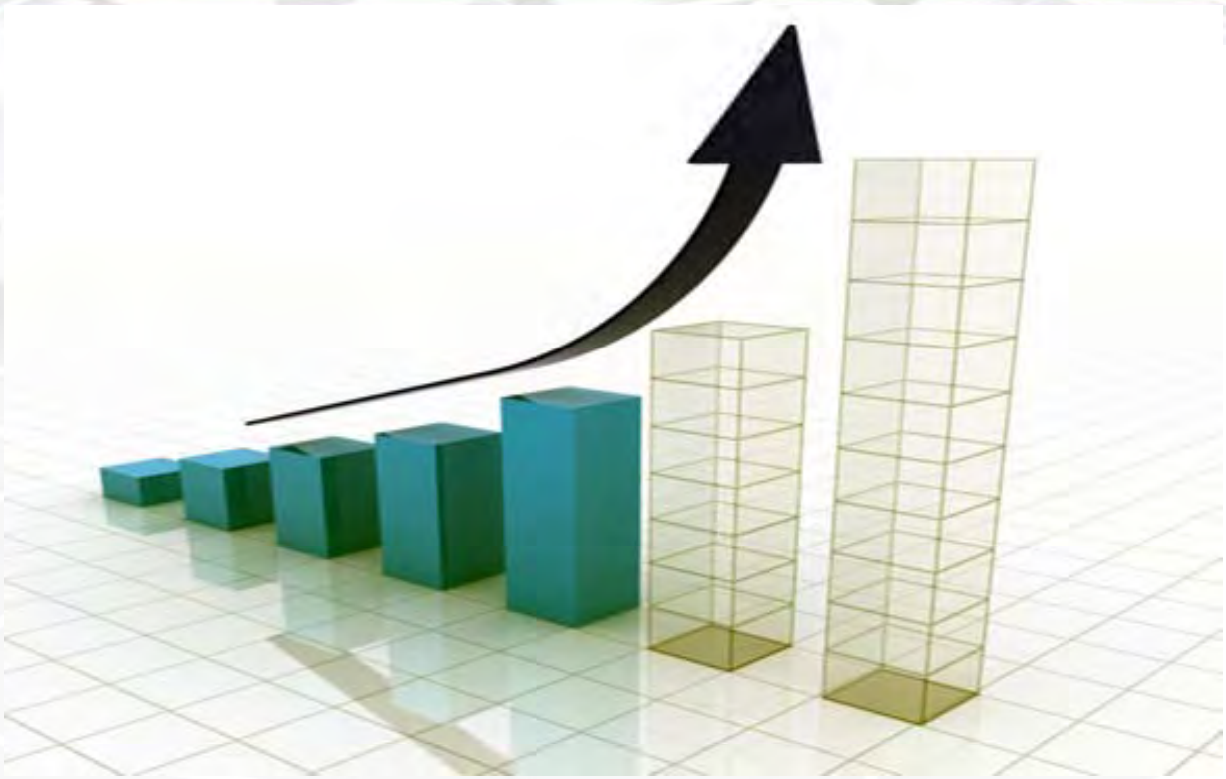

Data on network

Data to process

Computer speed

- CPU
- Bus
- Disk

GEMFIRE®
the enterprise data fabric

Decrease in:

GEMSTONE

- Decision time
- Time to build/deploy systems
- Time to adapt systems

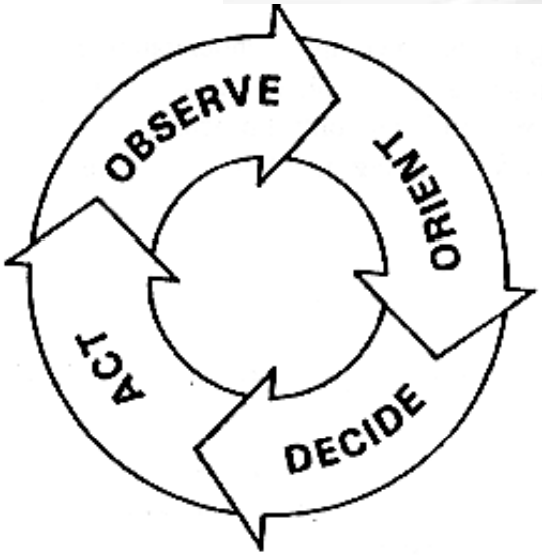

The command and control process:
The OODA loop.

GEMFIRE
the enterprise data fabric

Operational challenges…at the edge.

GEMSTONE®

- Power
- Space
- Bandwidth
- Reliability
- Availability
- Scalability

GEMFIRE®
the enterprise data fabric

Agenda

GEMSTONE®

- Key Trends
- The Challenge
- Enterprise Data Fabrics
- Use Cases

GEMFIRE®
the enterprise data fabric

Demands on IT

- Systems that keep pace with
-

ctical
-
-

IRE®
the enterprise
data fabric

Current mission challenges and technical pressures require new approaches

GEMSTONE

Deployed Data management technologies are not keeping pace with

- data volumes,
- compression of decision windows, and
- extension of the edge...

WHY?

- RDBMS are based on 1970's architecture based
- Speed of Light vs Speed Sound
- Data synchronization is not addressed by core technology
- Active / live in motion data vs. sleeping - poll driven data

GEMFIRE
the enterprise data fabric

Scale-Out Software Strategies

From Tier-Based Architectures to Partitioned Data Fabrics…

Partitioned data fabrics collocate the data and processing

Server-centric tier-based architectures separate the Data and the Processing

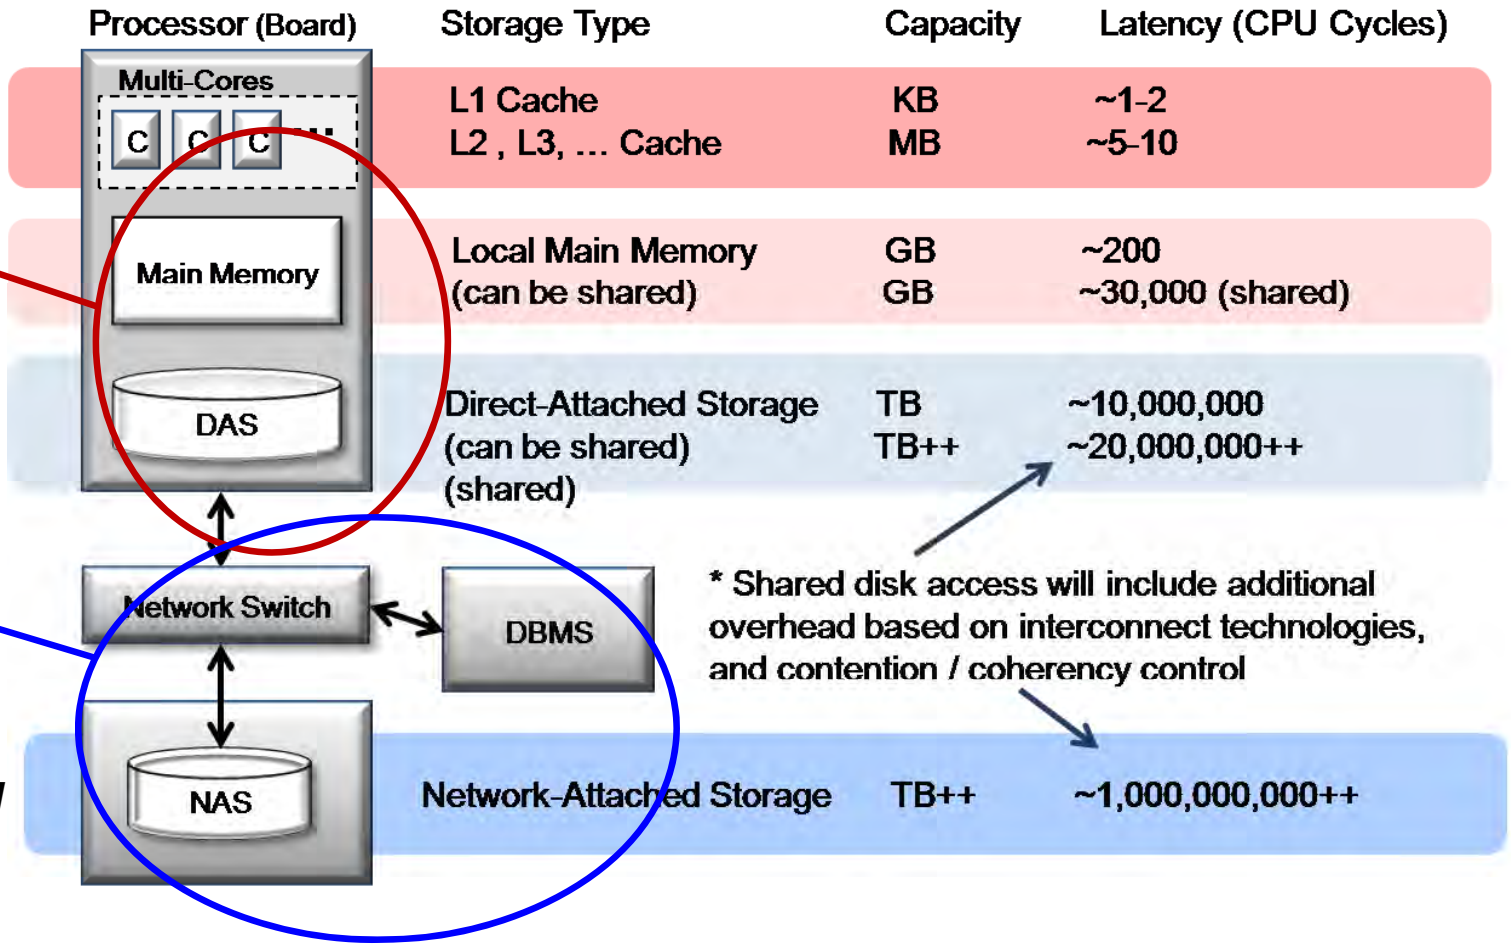

Traditional tier-based data strategies limit concurrency (and throughput) by creating contention, coherency, and latency effects

Public Release: 10-0861. Distribution Unlimited

© 2010 The MITRE Corporation. All rights reserved.

Current Mission Challenges and Technical Pressures require new approaches

GEMSTONE

Currently messaging needs are met by technology that is orthogonal to the "database“ -

- Requires fat pipes
- Separate processes
- Not dynamic
- Doesn’t know anything about the data
- Relies on disk...

GEMFIRE
the enterprise data fabric

Agenda

GEMSTONE®

- Key Trends
- The Challenge
- Enterprise Data Fabrics
- Use Cases

GEMFIRE®
the enterprise data fabric

So what is the ideal solution

GEMSTONE®

- 'speed of light' infinitely scalable data management with no bottle necks or single points of failure
- Distributed (WAN/LAN) with data synch
  - support for disconnected operations and low bandwidth environments
- Dynamic Membership (systems come and go)
- Built-in in-process messaging
- Query (like a database)
- Subscribe to Changes using Query
- Push processing to the data (Map-Reduce, Scatter-Gather)
- Secure
- Easy to insert into a system

GEMFIRE®
the enterprise data fabric

What is an Enterprise Data Fabric (EDF)?

GEMSTONE®

Distributed Cache +

- In Memory
- Distributed across a Cluster or Grid
- Synchronized
- Controlled Data Placement

Database

- Storage
- Persistence
- Transactions
- Queries
- High Availability
- Load Balancing
- Data Replication
- L1 Caching

+ Messaging System

- Data Distribution
- Event Propagation
- Guaranteed Delivery

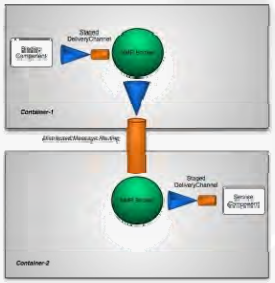

+ Low Latency Service Bus

- System Integration
- Data Transformation
- Service Loose Coupling

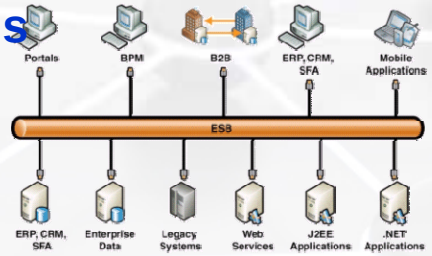

+ Grid Controller

- Task Decomposition
- Distributed Task Assignment
- Map-Reduce, Scatter-Gather
- Result Summarization

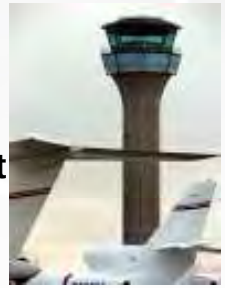

+ *Complex Event Processor*

- Fine Grained Event Detection
- Frame work for Real-time Analysis
- Enables Event Driven Architectures

An EDF uses select features from all of these products and combines them into a low-latency, linearly scalable, memory-based data fabric

GEMFIRE®
the enterprise data fabric

Why YAPOM ? Yet Another Piece of Middleware?

GEMSTONE

- **Execution Excellence**
  - Make your applications run 4 to 40 times faster
  - Ingest / Digest / Distribute vast amounts of data with extremely low latency
- **Higher ROI of IT investments & Lower Cost Per Transaction**
  - A 4x Performance Increase is like getting **3 free computers** that take **no space or electricity**
- **Enables applications to survive network outages and distressed networks**
- **Data Awareness – especially across applications/systems**
- **Supports High Availability, Fault Tolerance, and Site Failover with zero additional design/development costs**

Agenda

GEMSTONE®

- Key Trends
- The Challenge
- Enterprise Data Fabrics
- Use Cases

GEMFIRE®
the enterprise data fabric

Use Cases

GEMSTONE®

- Multistage Data Processing
  - Increase performance, reduce costs
- Common Data Backbone
  - Scale economically and incrementally, reduce costs
  - Break data stovepipes (not process),
  - Detection and propagation of events, simple and reliable
- Barrier Event Detection
  - Performance, Data Awareness
- Multisite Data Management
  - Share data globally, built-in COOP and HA
- Real Time Situational Awareness
  - Fine grained data subscriptions, event detection, DIL/disconnected ops
- Data Awareness
  - High Performance SOA, Break data stovepipes, Data driven eventing
- Joint Space Operations Center Prototype (and Wall Street Risk Management)
  - High Speed Data Ingest
- US Intelligence Agency Detection & Alerting Solution
  - Low Latency Service Bus

GEMFIRE®
the enterprise data fabric

Customer Use:
Sequential Transaction Processing

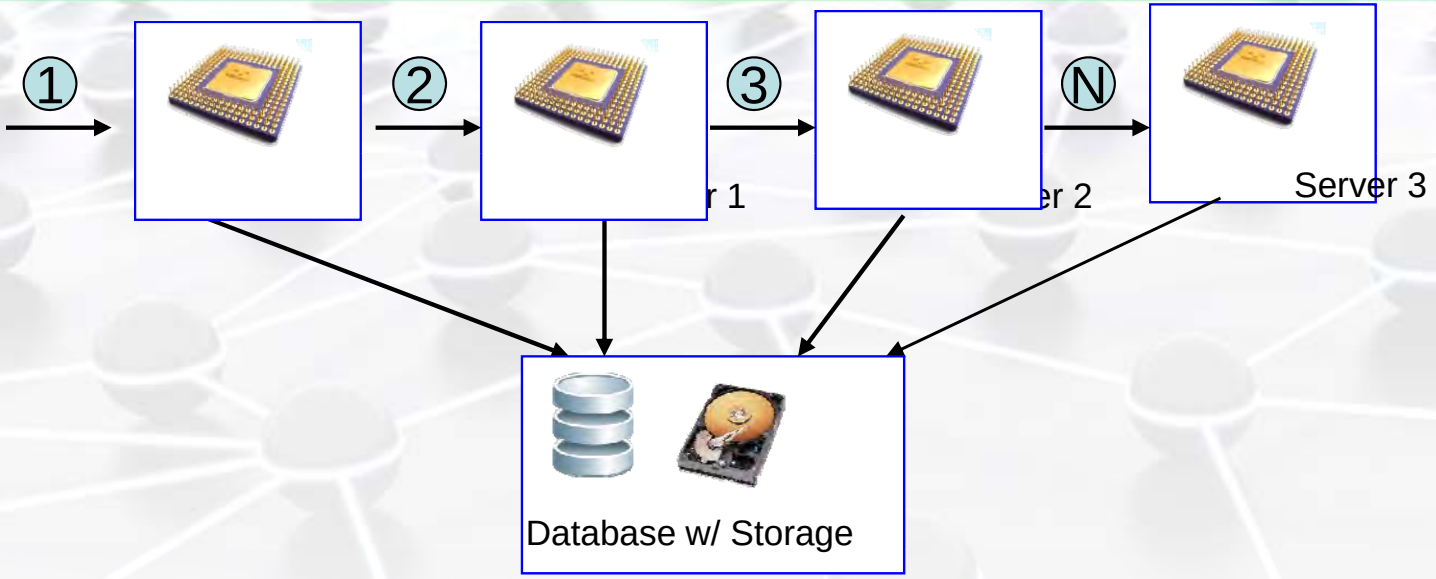

Today's "traditional" multi-stage transaction processing architecture

Using RDBMS's
Message Buses
Or Both

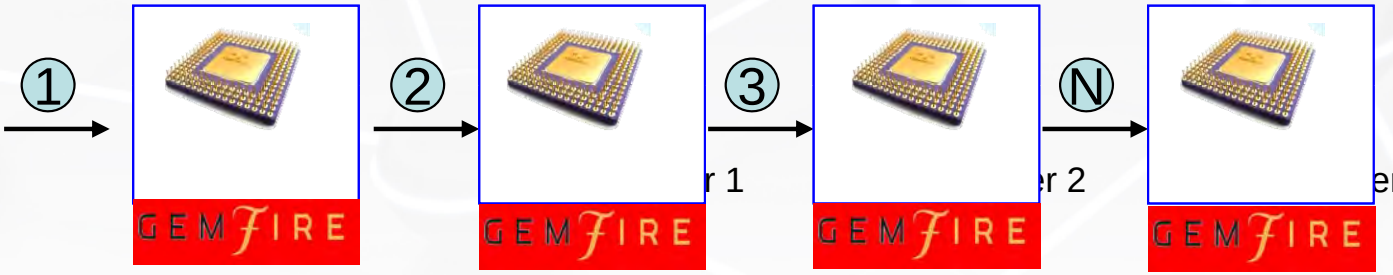

Customers such as JPMC, UBS, Nomura, Merrill, have seen a 4 to 40 times processing speed up & reduced end to end latency

Process more data
Reduce hardware
Add more calculations

GEMFIRE
Optional GemFire Server

GEMFIRE
the enterprise data fabric

Data Management Solution - Common Data Backbone

GEMSTONE

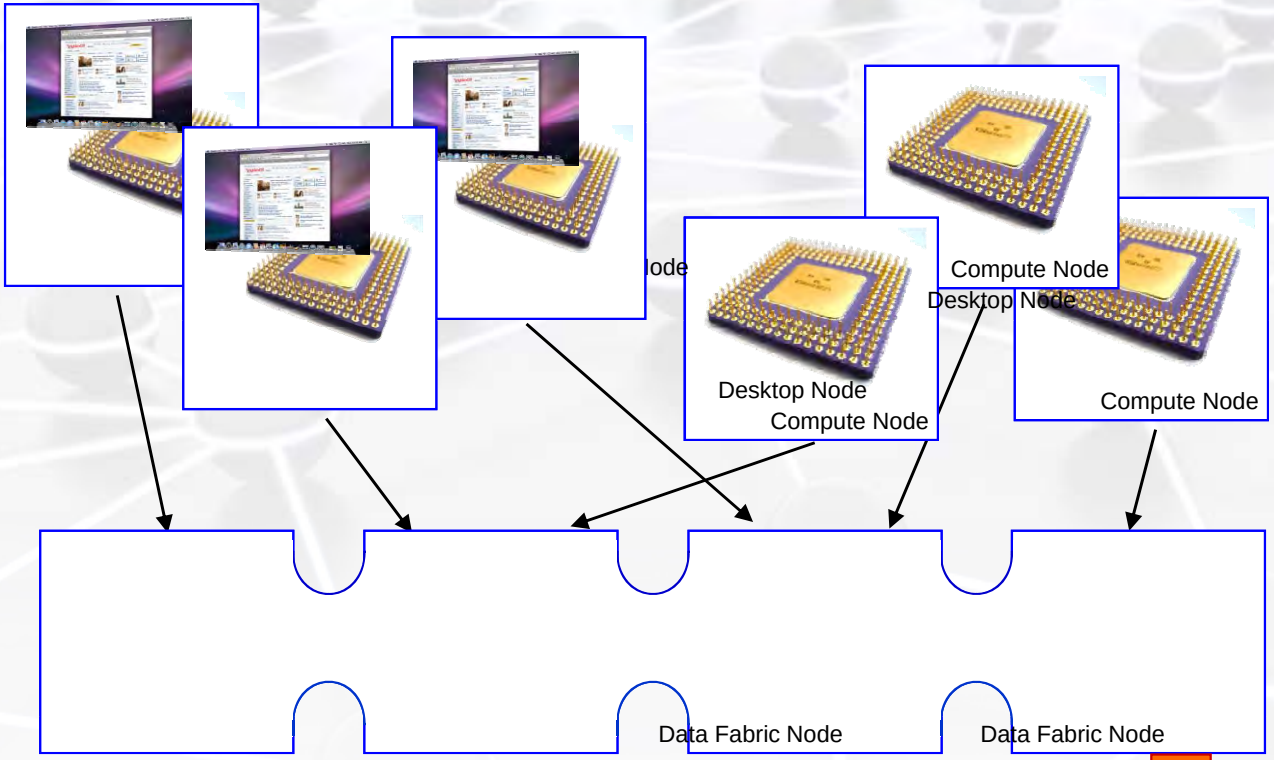

- Shared Resources
- Predictable Latency
- Scalable
- Consistent Data Models
- Reports Reconcile
- Common Functions/Code
- Real Time
- Events Propagated

When the Data Fabric Holds objects:

- O-R Mapping Churn is avoided (saving CPU & I/O)
- Common Java code can execute in the fabric

When the Data Fabric Holds Relational Tables:

- Existing SQL Applications can plug in more easily
- Use Distributed Functions for common code

Asynchronous,
Conflated,
Batched Updates

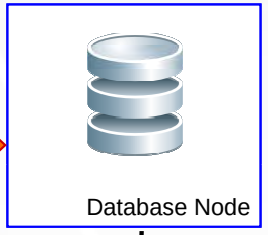

Archival, OLAP &
Regulatory RDBMS

Storage Device

GEMFIRE
the enterprise data fabric

Common Data Backbone

GEMSTONE®

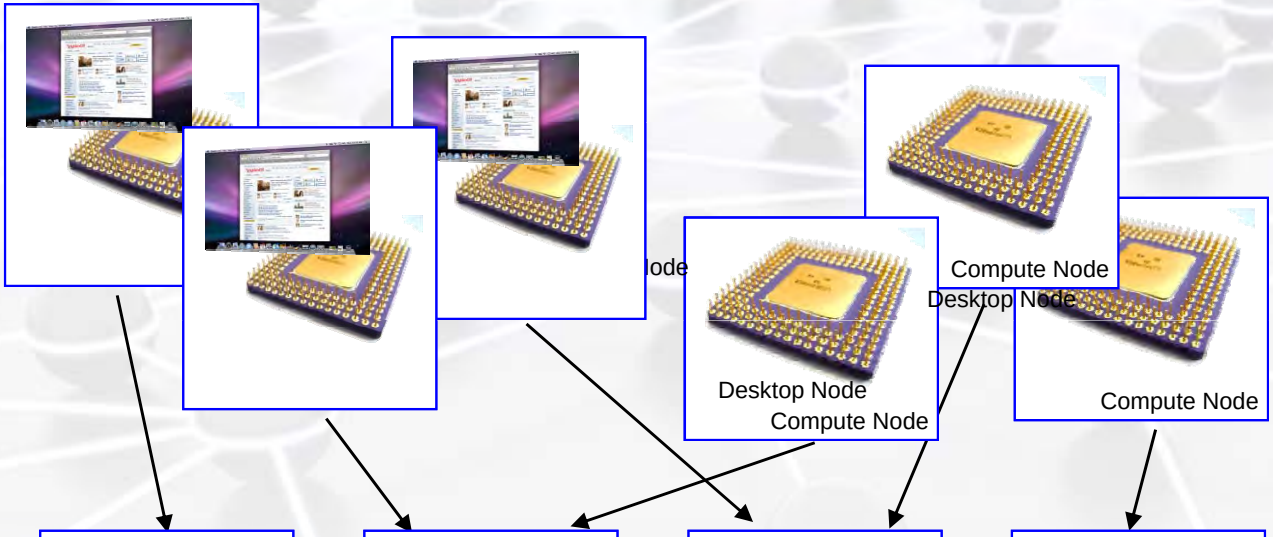

- Counterparty Descriptions
- Settlement Instructions
- Netting Agreements
- Instrument Data
- Rating Information

Data Fabric Node

Data Fabric Node

Replicated Regions

Data Fabric Node

Data Fabric Node

Asynchronous,
Conflated,
Batched Updates

Database Node

Archival, OLAP &
Regulatory RDBMS

Storage Device

GEMFIRE®
the enterprise
data fabric

Common Data Backbone

GEMSTONE®

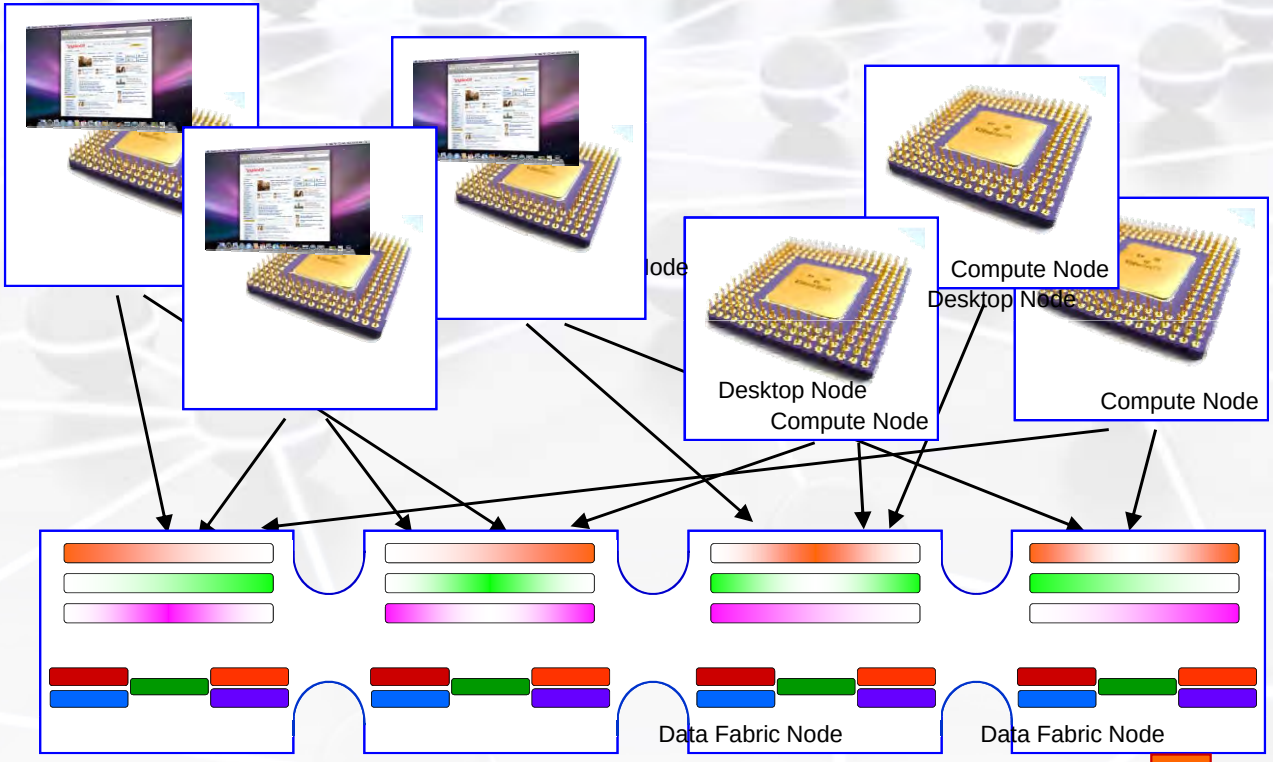

Position Data

Trade Data

Market Data

Active Data

GEMFIRE®
the enterprise
data fabric

Common Backbone = Information as a Service

GEMSTONE

GEMFIRE real-life use case: Common Application Backbone

- Shared Infrastructure
  - Multiple business critical functions can share common infrastructure, common data, common calculations
  - Reduces data-center cost/complexity
- Continuous availability
  - All data is always available as and when needed.
  - No single point of failure
  - Obviates the need for third party H/A clustering software
- Common analytics for all calculations
  - Intra day trader risk and end of day Flash P&L, and close-of-day risk batch all use the exact same inputs and calculations so results are more precise and easier to reconcile

GEMFIRE
the enterprise data fabric

Compute Grids Deployed with GemFire

GEMSTONE

GEMFIRE real life use case: Major Bank

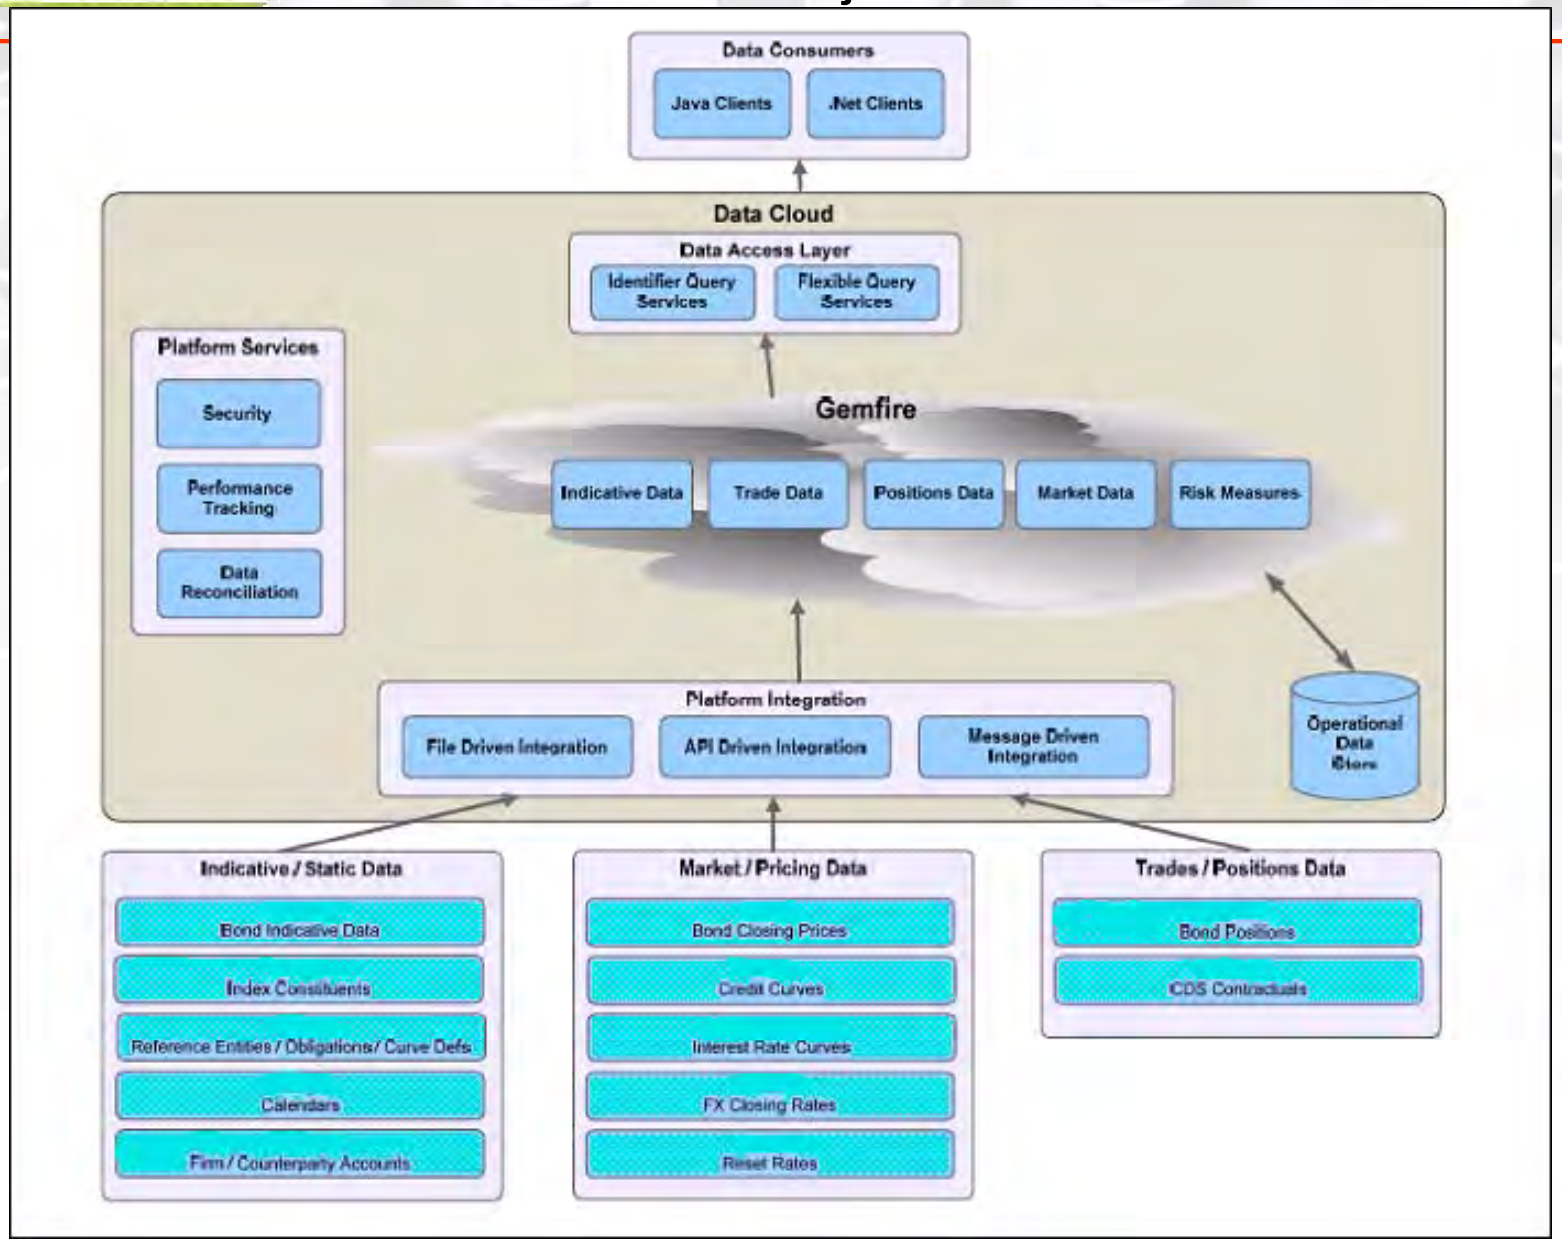

GEMFIRE
the enterprise data fabric

Barrier Event Detection

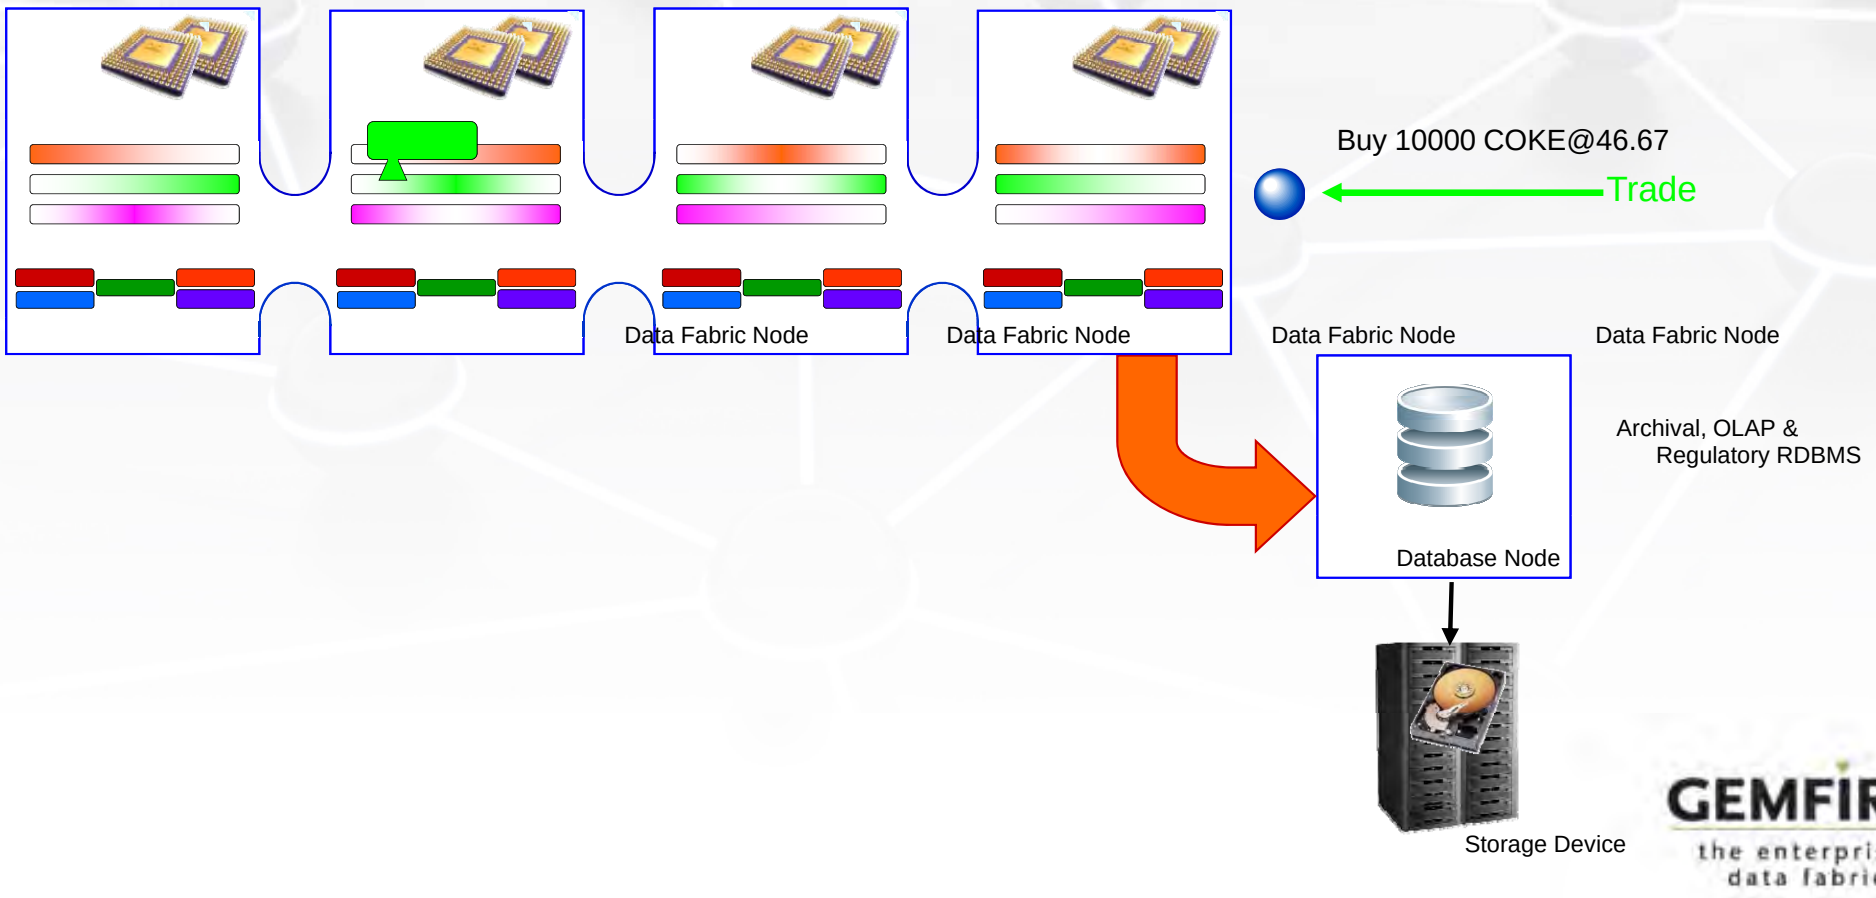

GEMFIRE
the enterprise data fabric

Barrier Event Detection

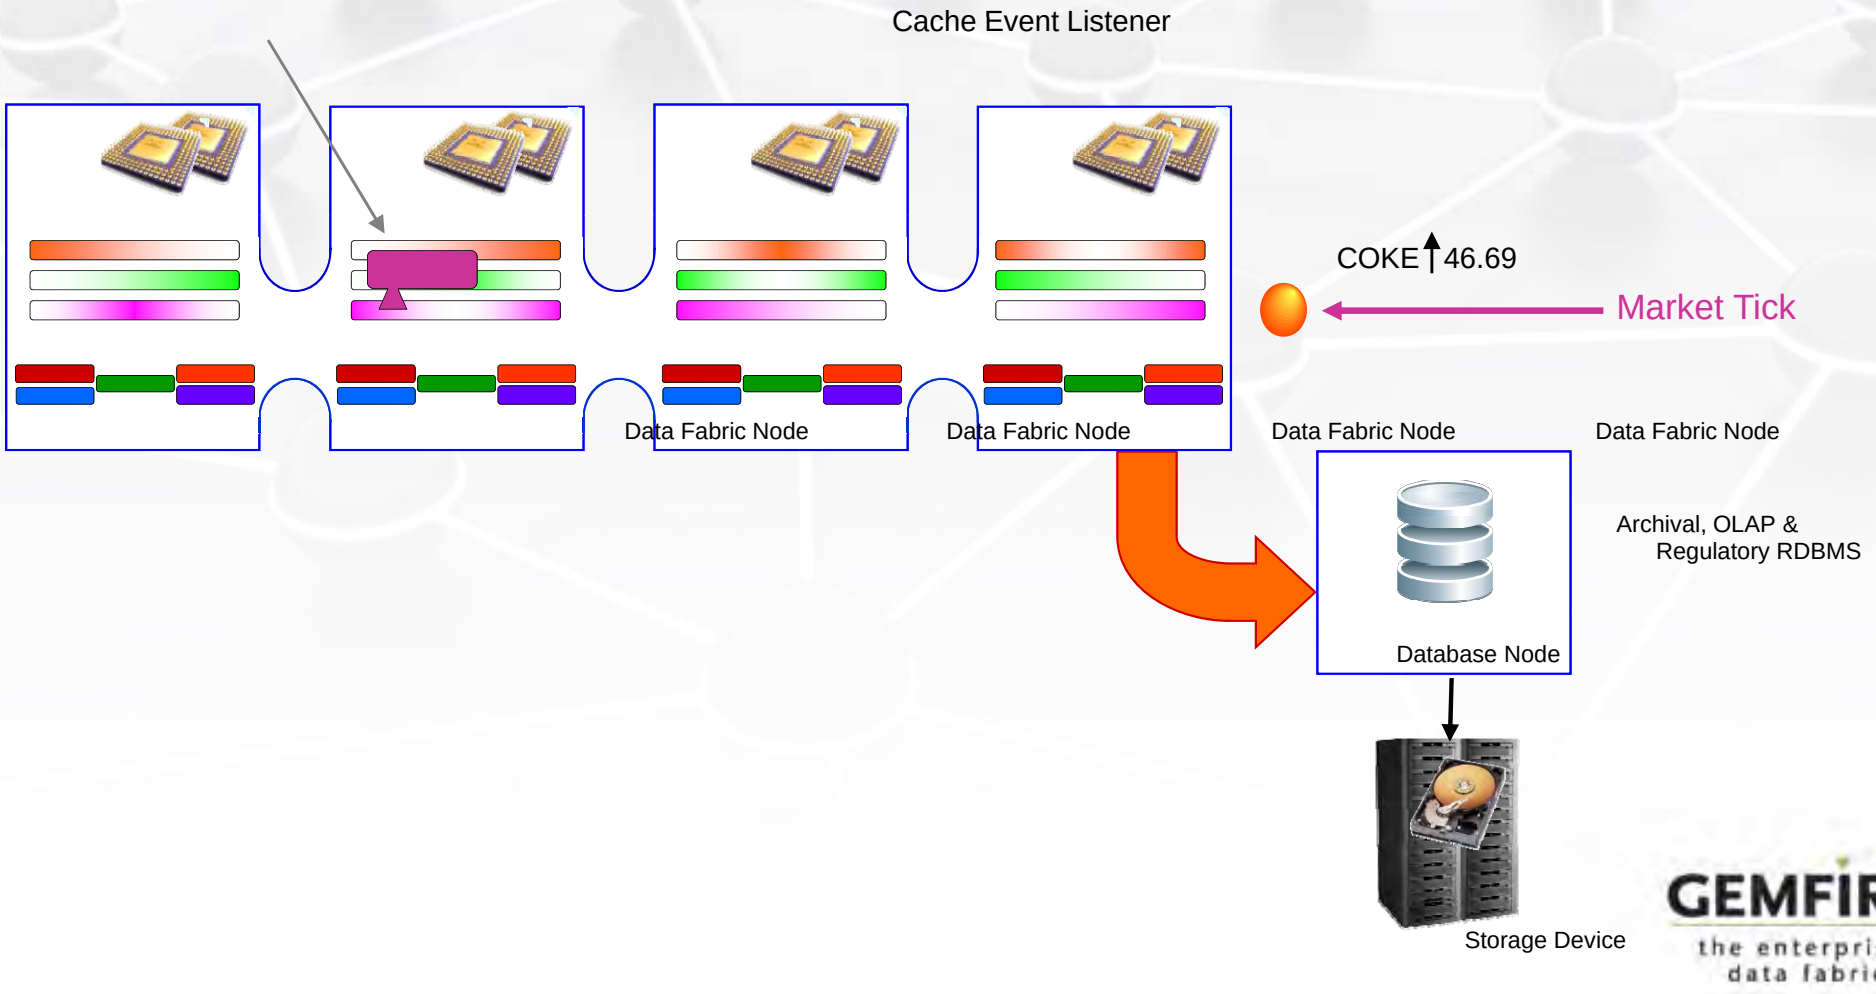

Barrier Event Detection

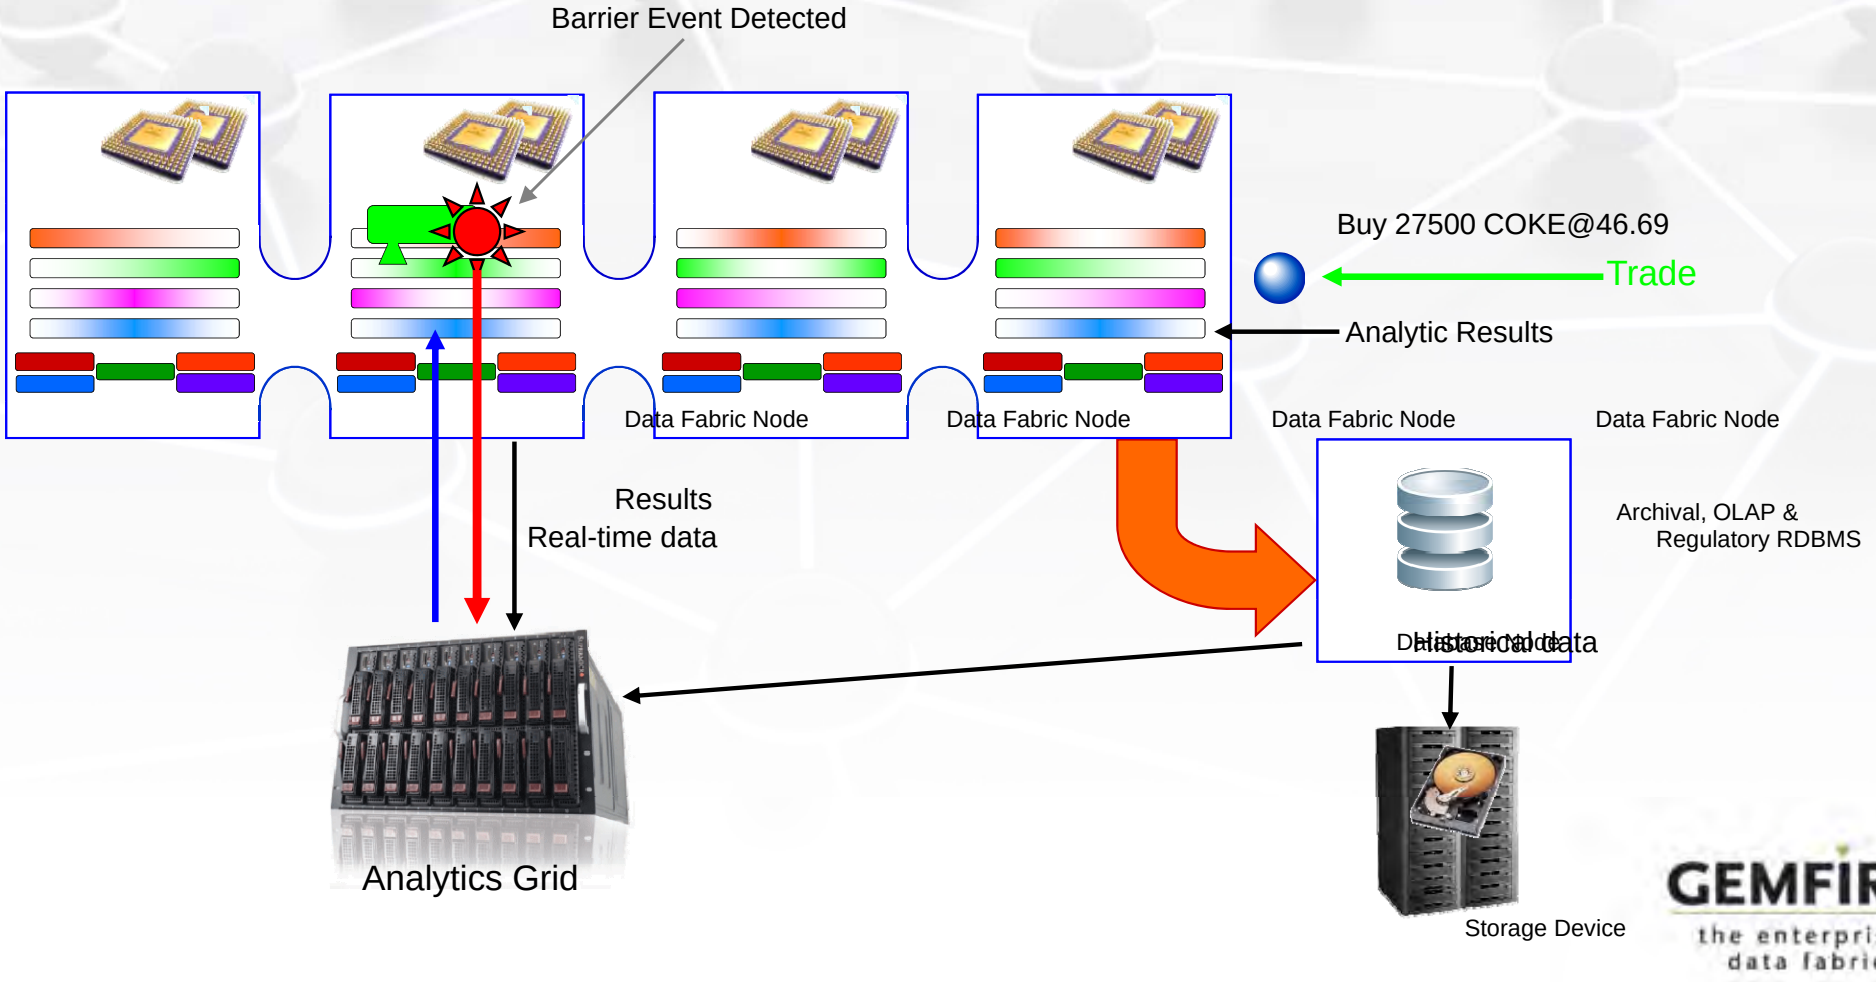

Barrier Event Detection

GEMSTONE

GEMFIRE real-life use case: Event-driven Risk and Pricing

Batch processing delays were costing money and limiting opportunities

JPMorgan

Running Pricing and Risk calculations in batch mode required special applications to be written that created a 'buffer' to ensure that regulatory obligations were not violated.

With GemFire:

Real-time event detection, easier integration of new products, and control workflow to minimize business risk and meet all regulatory requirements

8 CPUs re-price 6000 complex and exotic instruments in under 3 minutes

No longer need to tie up so much of our working capital to "run the bank"

The Banker
Innovation in data management technology award 2009

GEMFIRE
the enterprise data fabric

Global Book

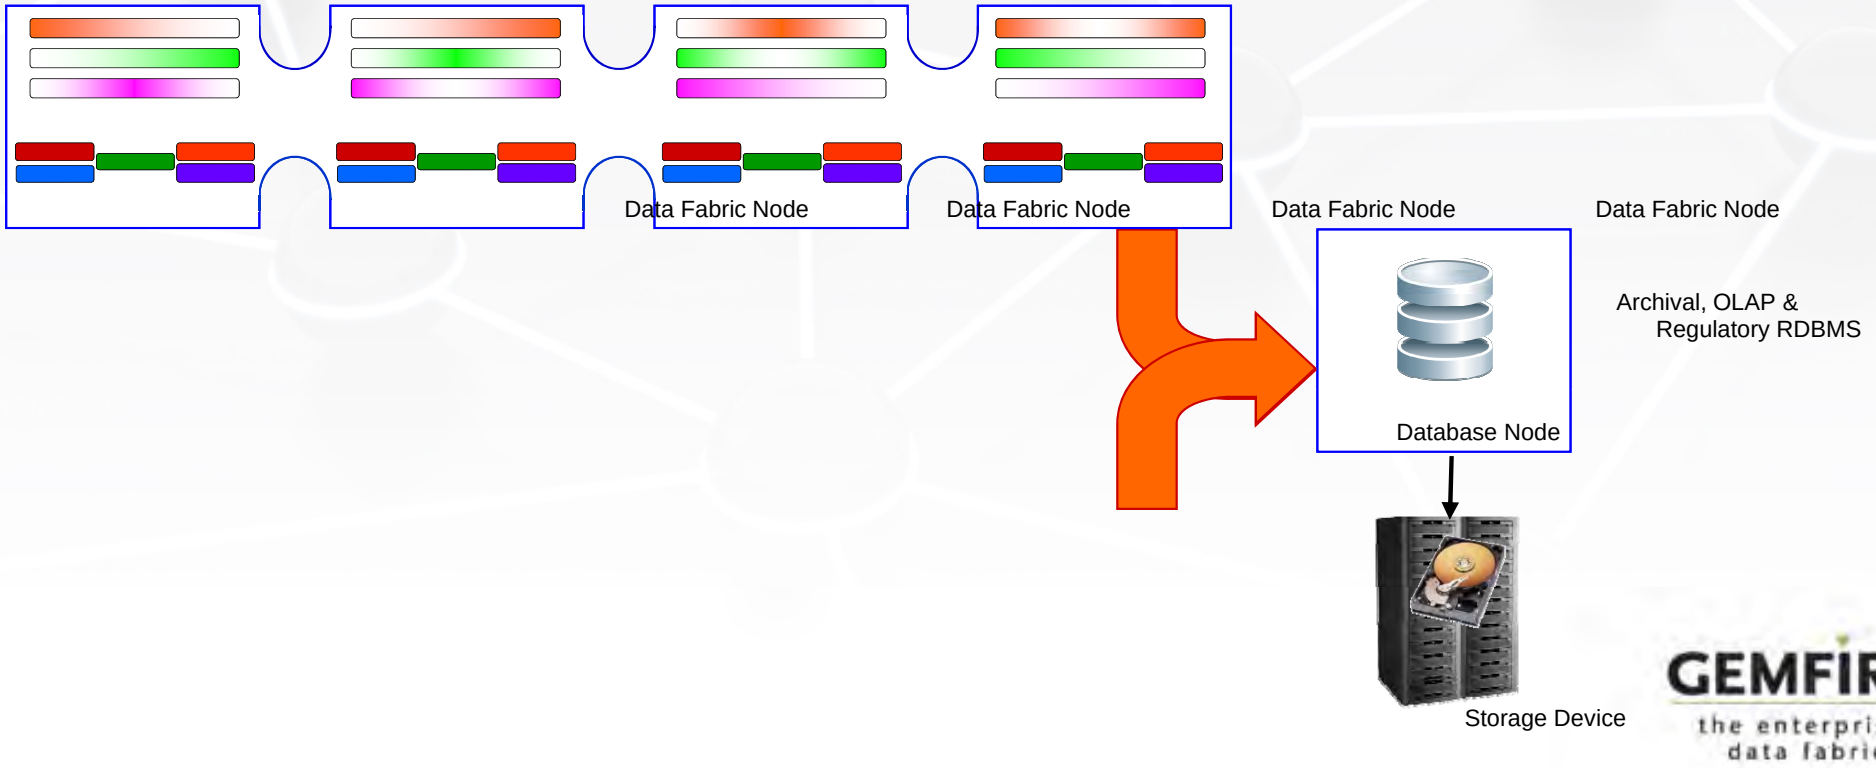
Data Fabric Node
Data Fabric Node
Data Fabric Node
Data Fabric Node
Archival, OLAP &
Regulatory RDBMS
Database Node
Storage Device

Global Book

GEMSTONE

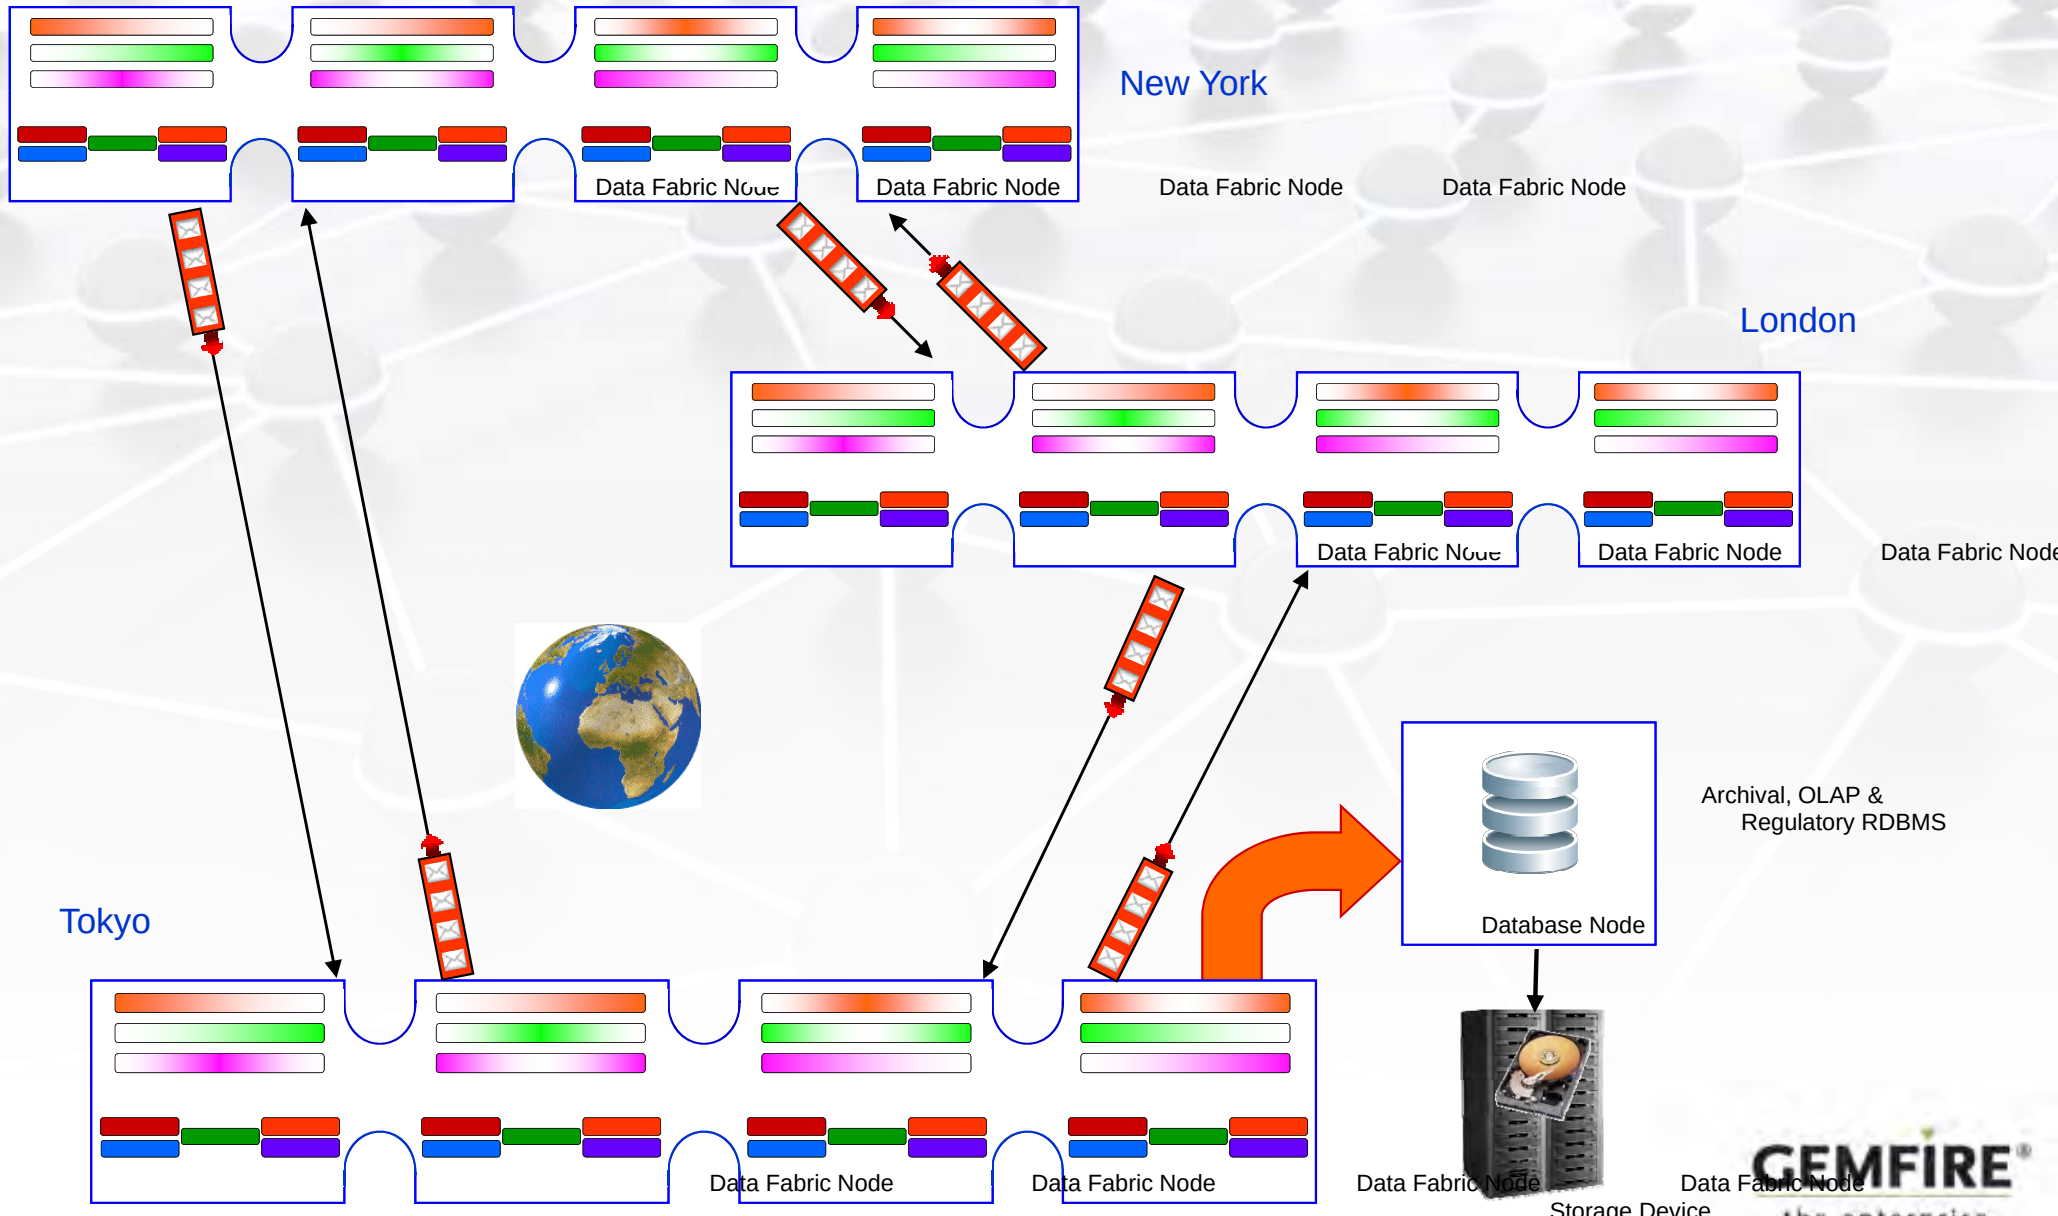

GCCS-J and GemFire Agile Client

GEMSTONE

User Interface with

Continuously updated

Fused Operational Picture

Client w/ GemFire

Client w/ GemFire

Continuous Queries

Client Initiated, Server Side Filter

Correlation Servers

Site 1

TMS-S Servers

Site 2

...

TMS-S Servers

Site N

70+ data feeds

GEMFIRE
the enterprise data fabric

Extending Real-time Data Analysis

GEMSTONE

Continuous Analytics that fuse multiple data sources

Local Query and Analysis

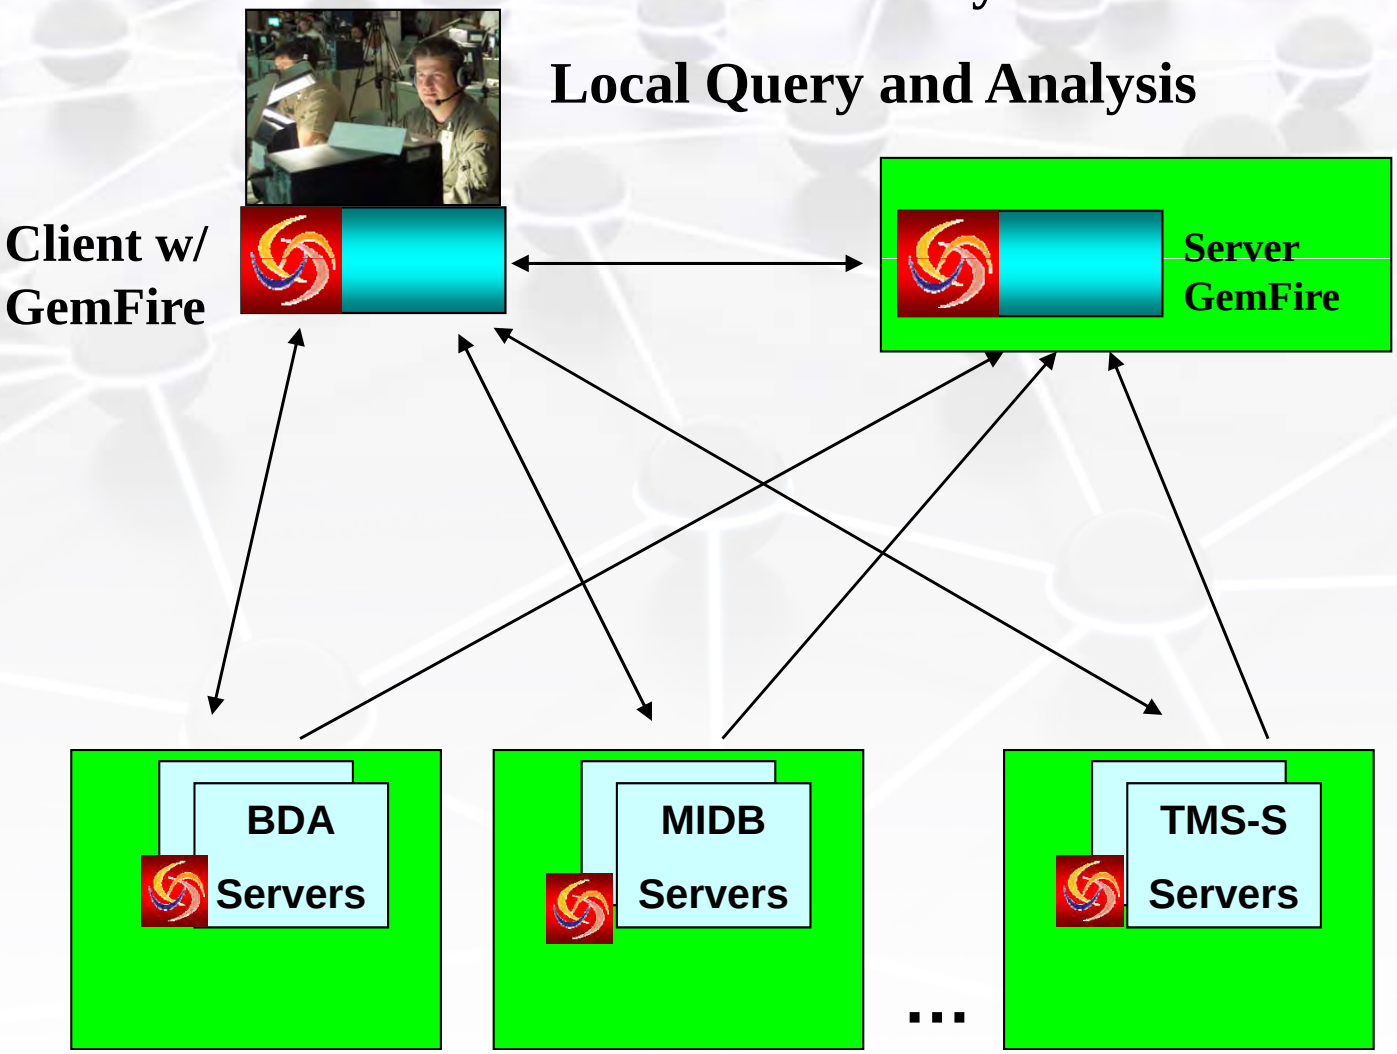

GEMFIRE
the enterprise
data fabric

GCCS-J: JOPES SSE, SORTS SSE and Critical Global Sites

NORTHROP GRUMMAN

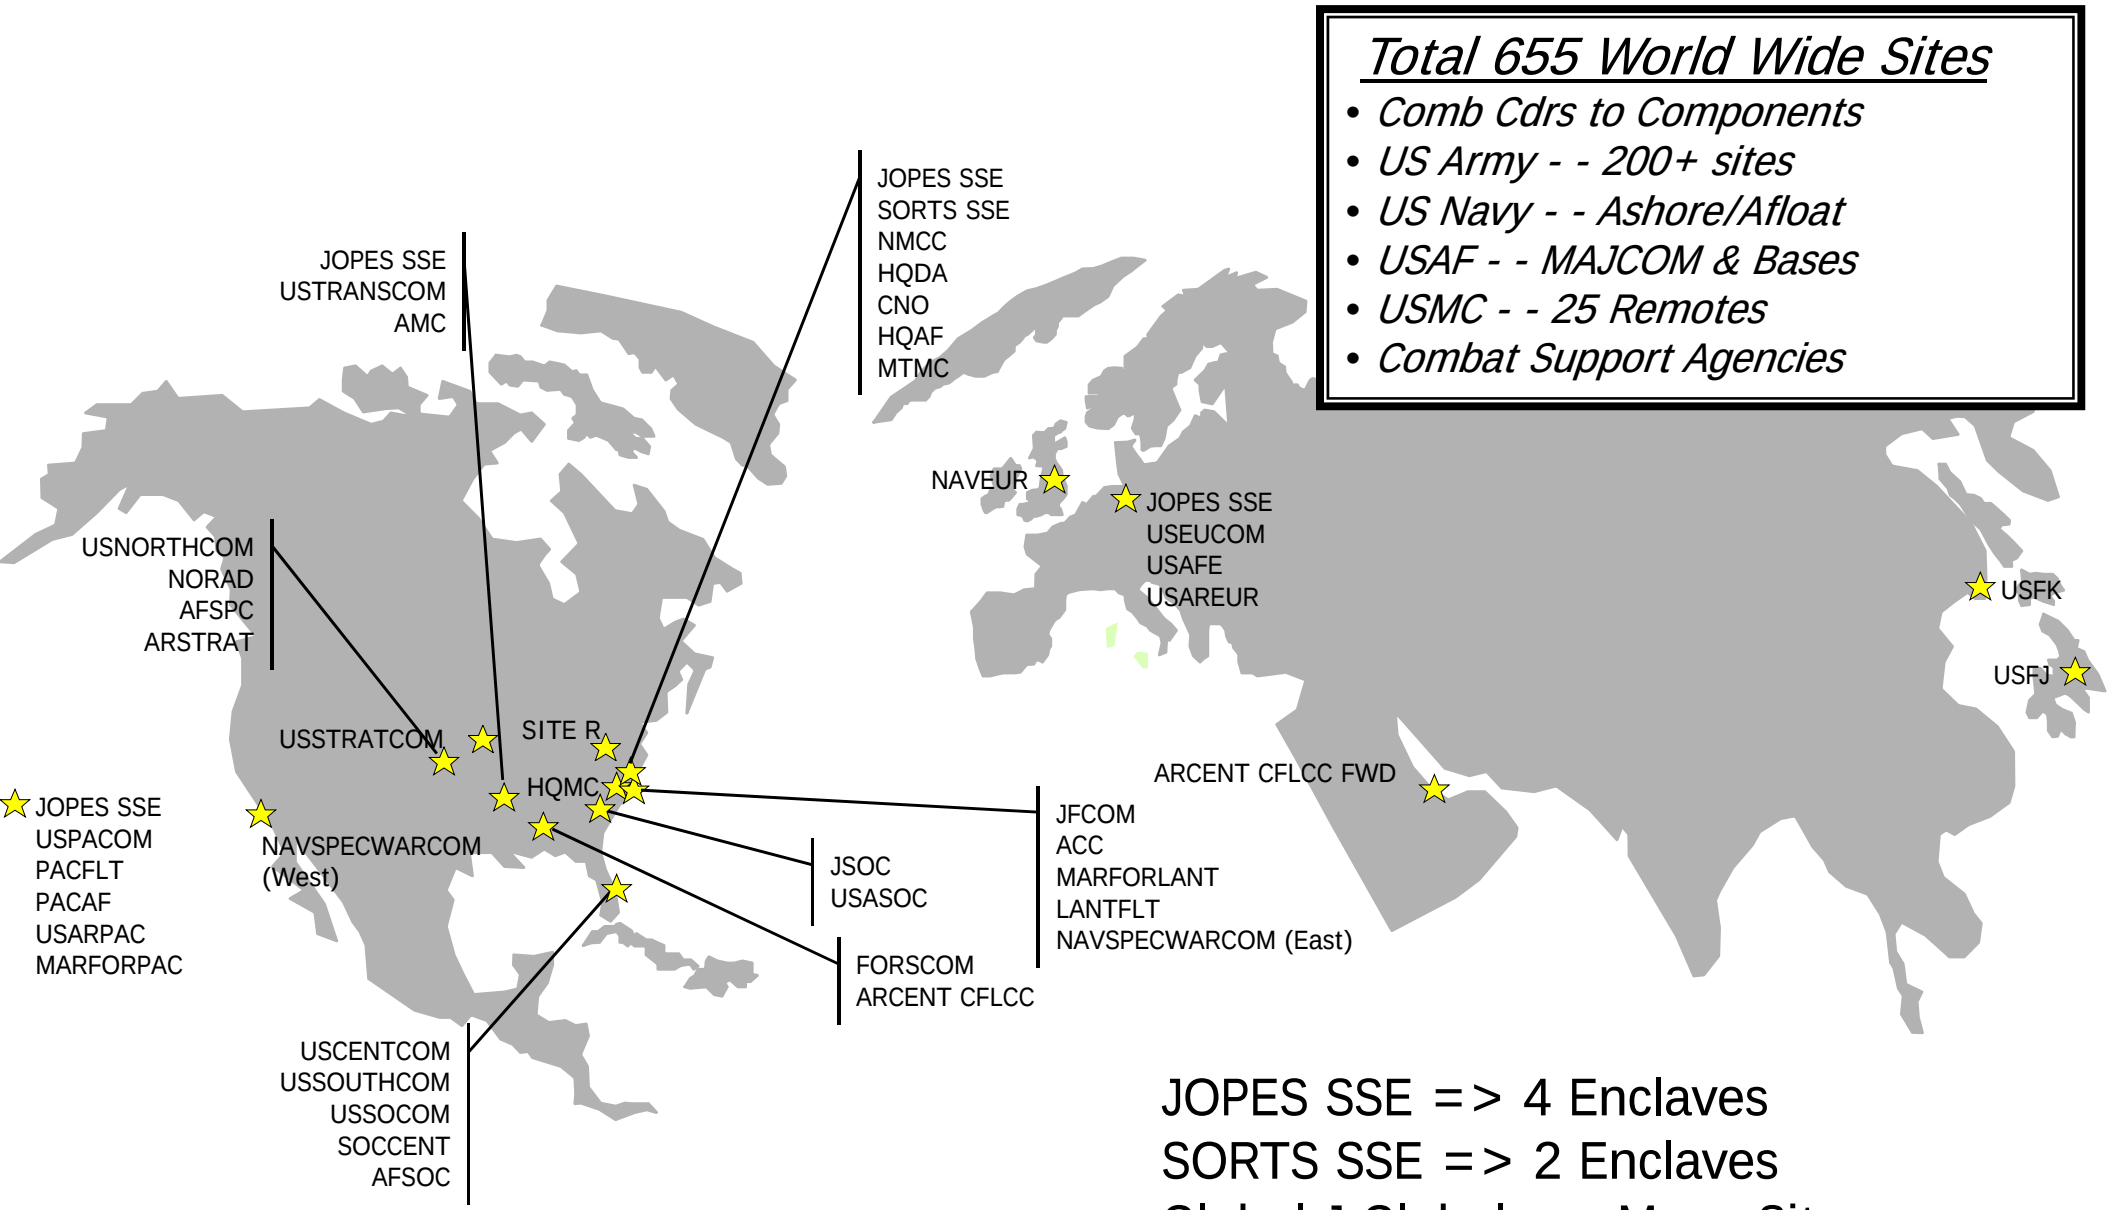

Total 655 World Wide Sites

- *Comb Cdrs to Components*
- *US Army - - 200+ sites*
- *US Navy - - Ashore/Afloat*
- *USAF - - MAJCOM & Bases*
- *USMC - - 25 Remotes*
- *Combat Support Agencies*

JOPES SSE => 4 Enclaves
SORTS SSE => 2 Enclaves
Global-J Global => Many Sites

NORTHROP GRUMMAN PROPRIETARY LEVEL I

Application Collaboration Today

GEMSTONE

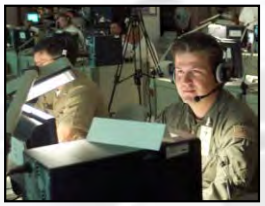

Mission Planning

Battle Damage Assessment

1. Mission planning starts to
- Destroy enemy SAM position
- Gets data from other sources
  -- BDA, MIDB, etc

3. Mission planner syncs with sources before finishing mission plan; sees changes
   Reworks mission to target 2

5. Mission executer deals with changes at run-time.

2. BDA update - SAM position is destroyed
   Mission planner does NOT see this !

4. BDA update – target 2 destroyed
   Mission planner does NOT see this !

GEMFIRE
the enterprise data fabric

Application Collaboration Tomorrow

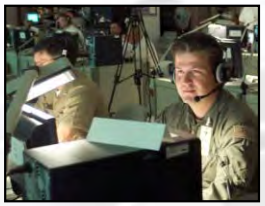

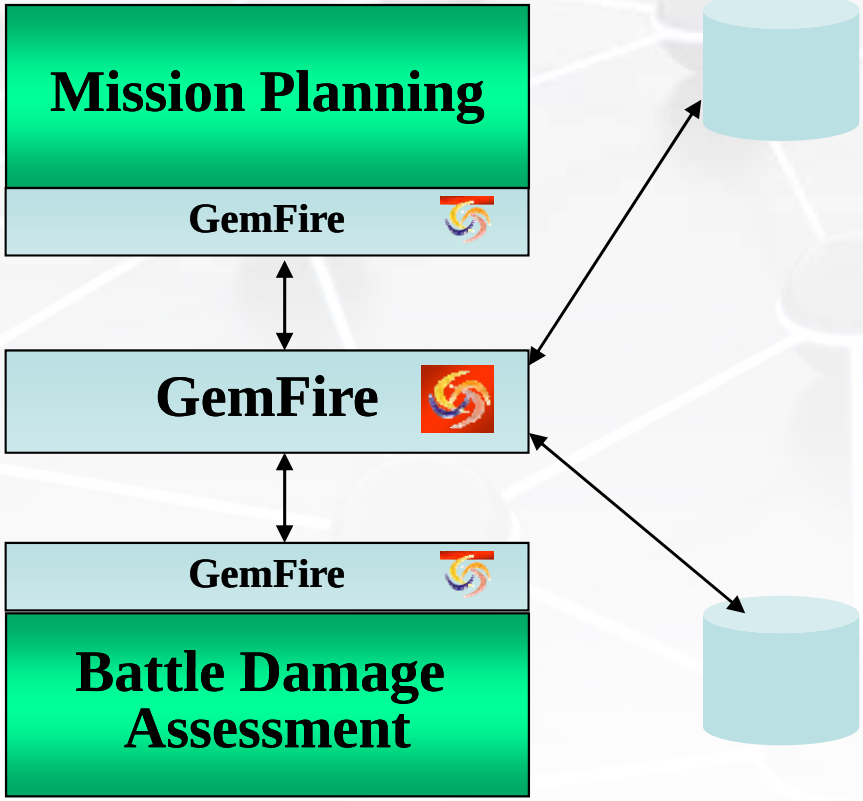

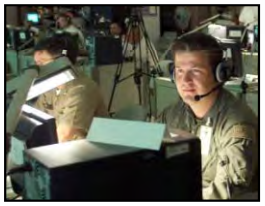

1. Mission planning starts to
   - Destroy enemy SAM position
   - Gets data from other sources
     -- BDA, MIDB, etc

3. Mission planner sees change
   Reworks mission to target 2

5. Mission executer is notified on way to target and is sent an alternate target

2. BDA update - SAM position is destroyed
   Mission planner gets change notification

4. BDA update after mission is started– target 2 destroyed
   Mission planner's application is notified

Risk Management

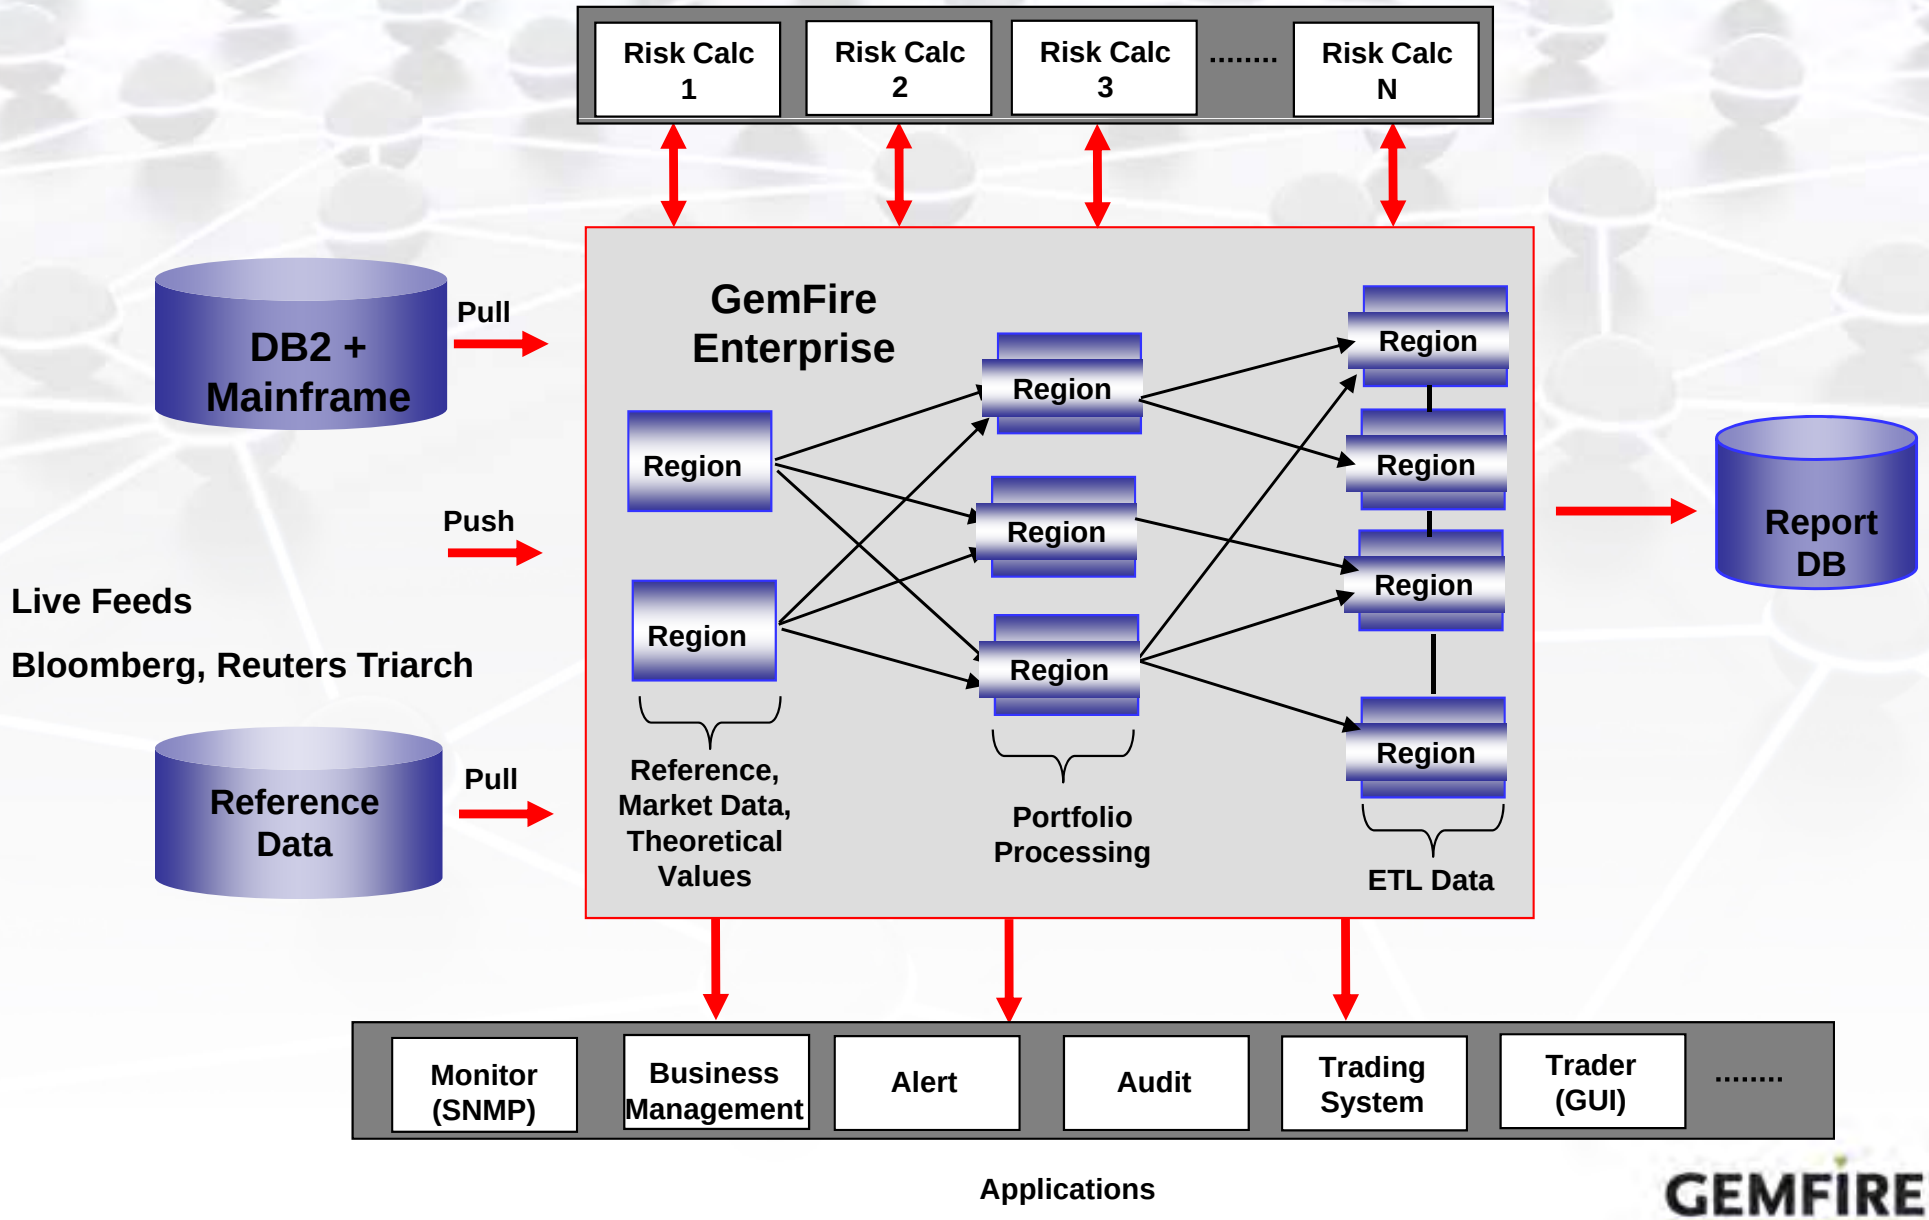

GEMFIRE
the enterprise data fabric

JSPOC

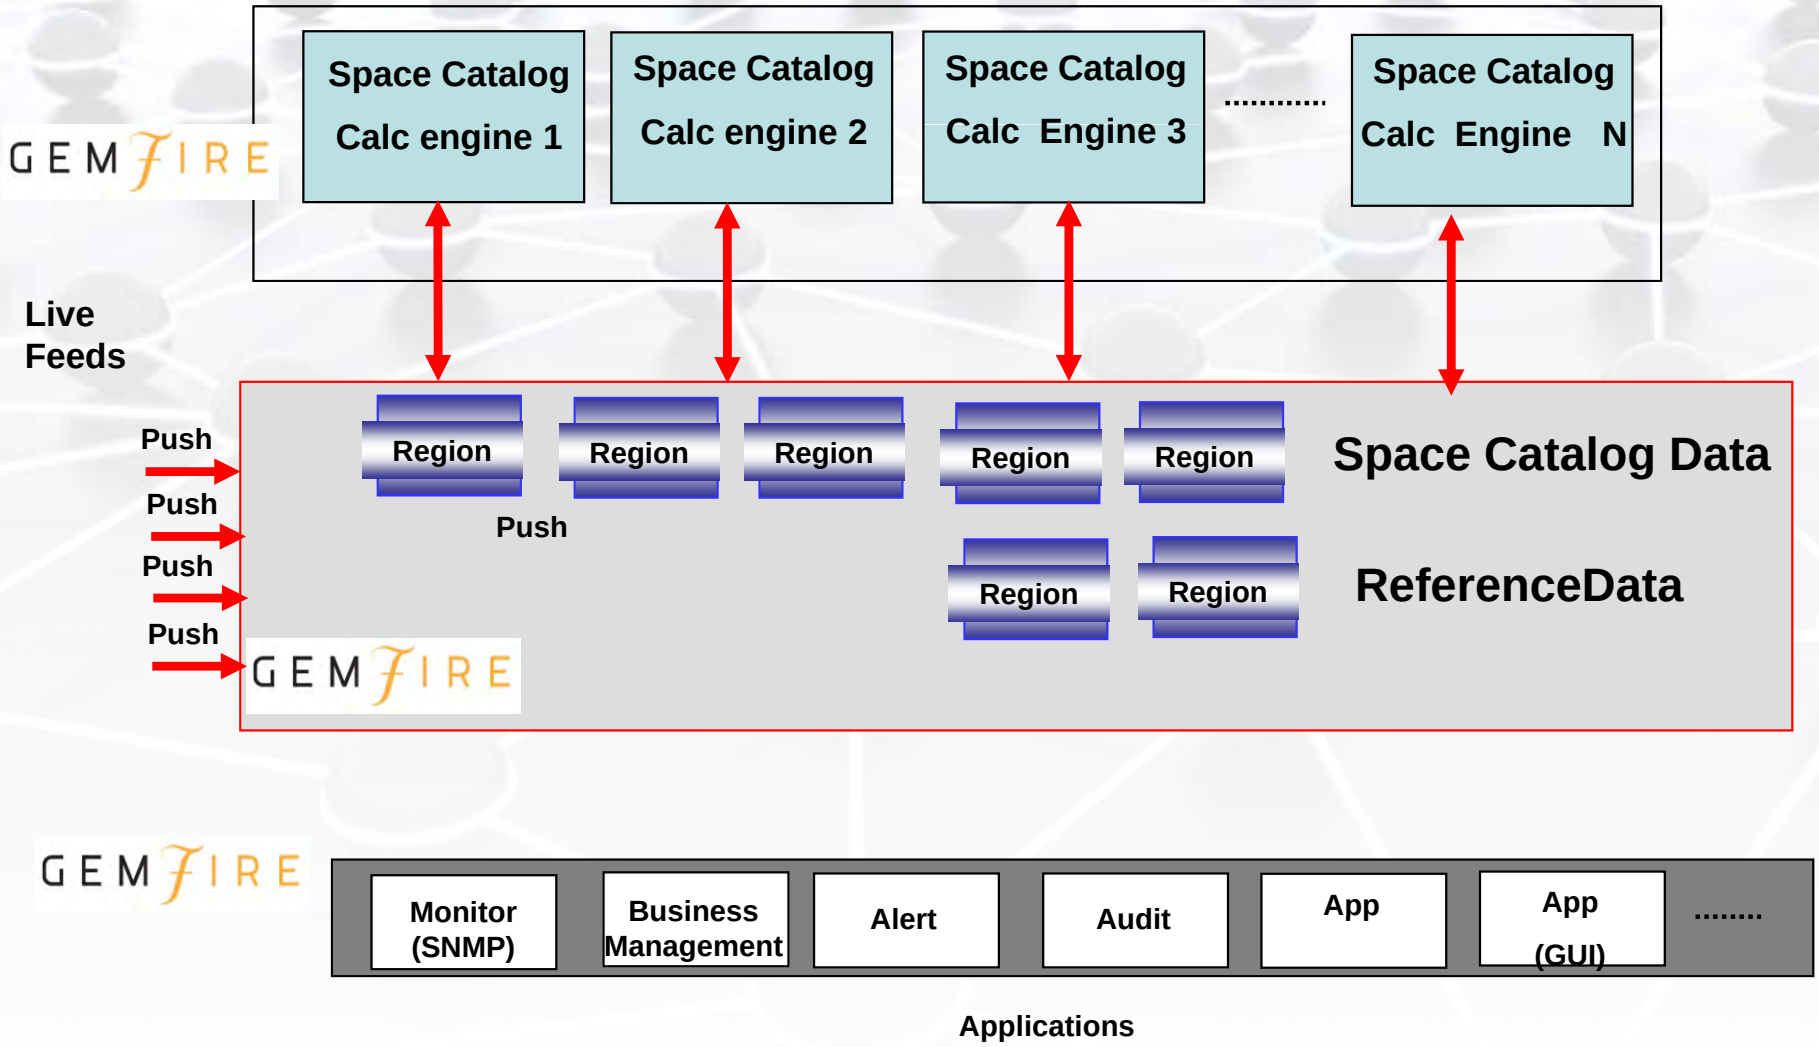

GEMFIRE
the enterprise
data fabric

EDF as High Speed SOA

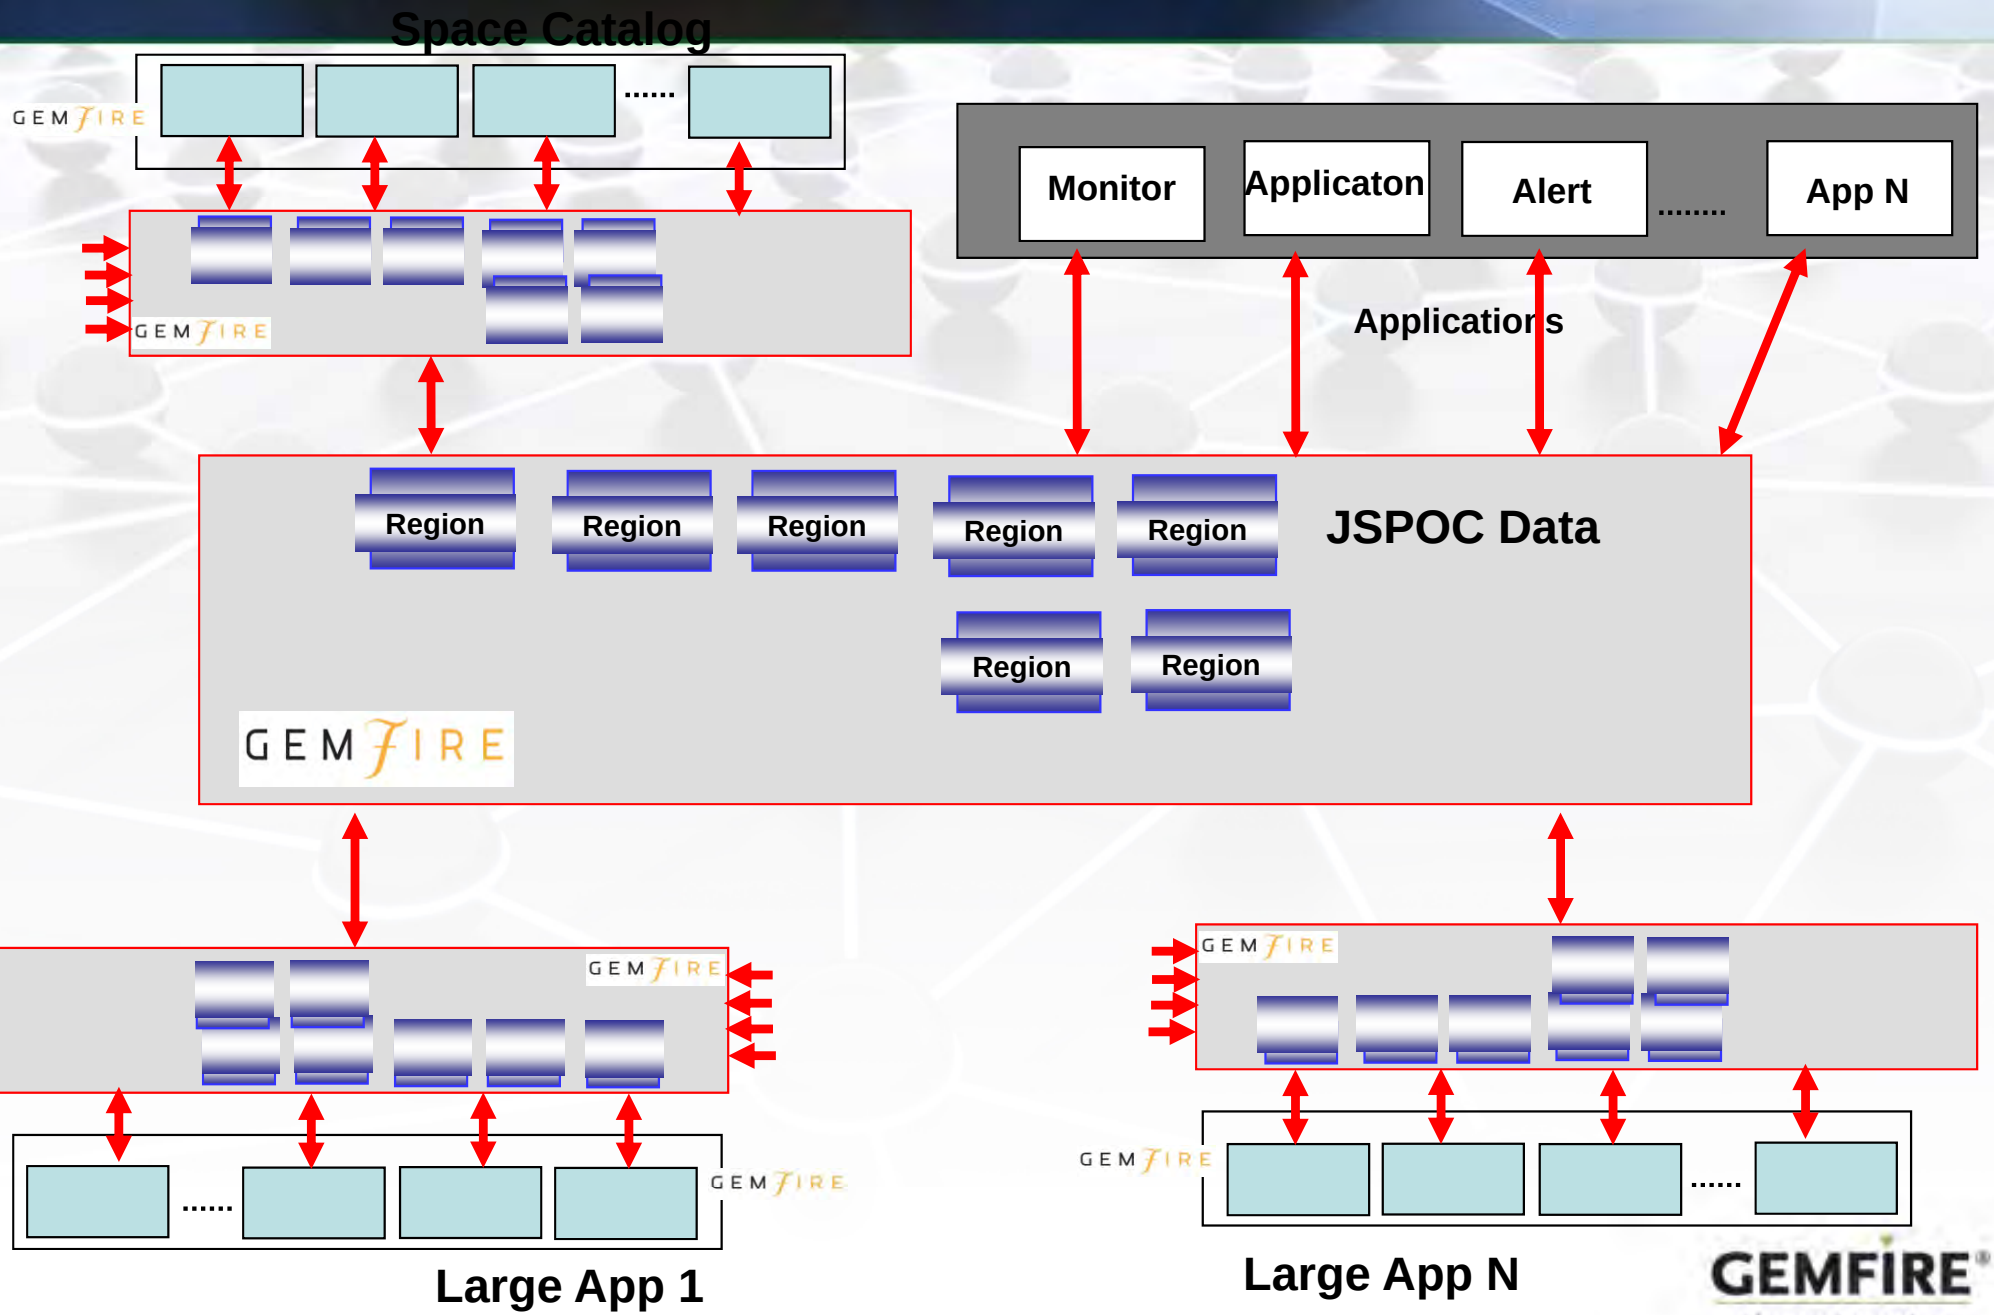

Detection & Alerting Solution Overview

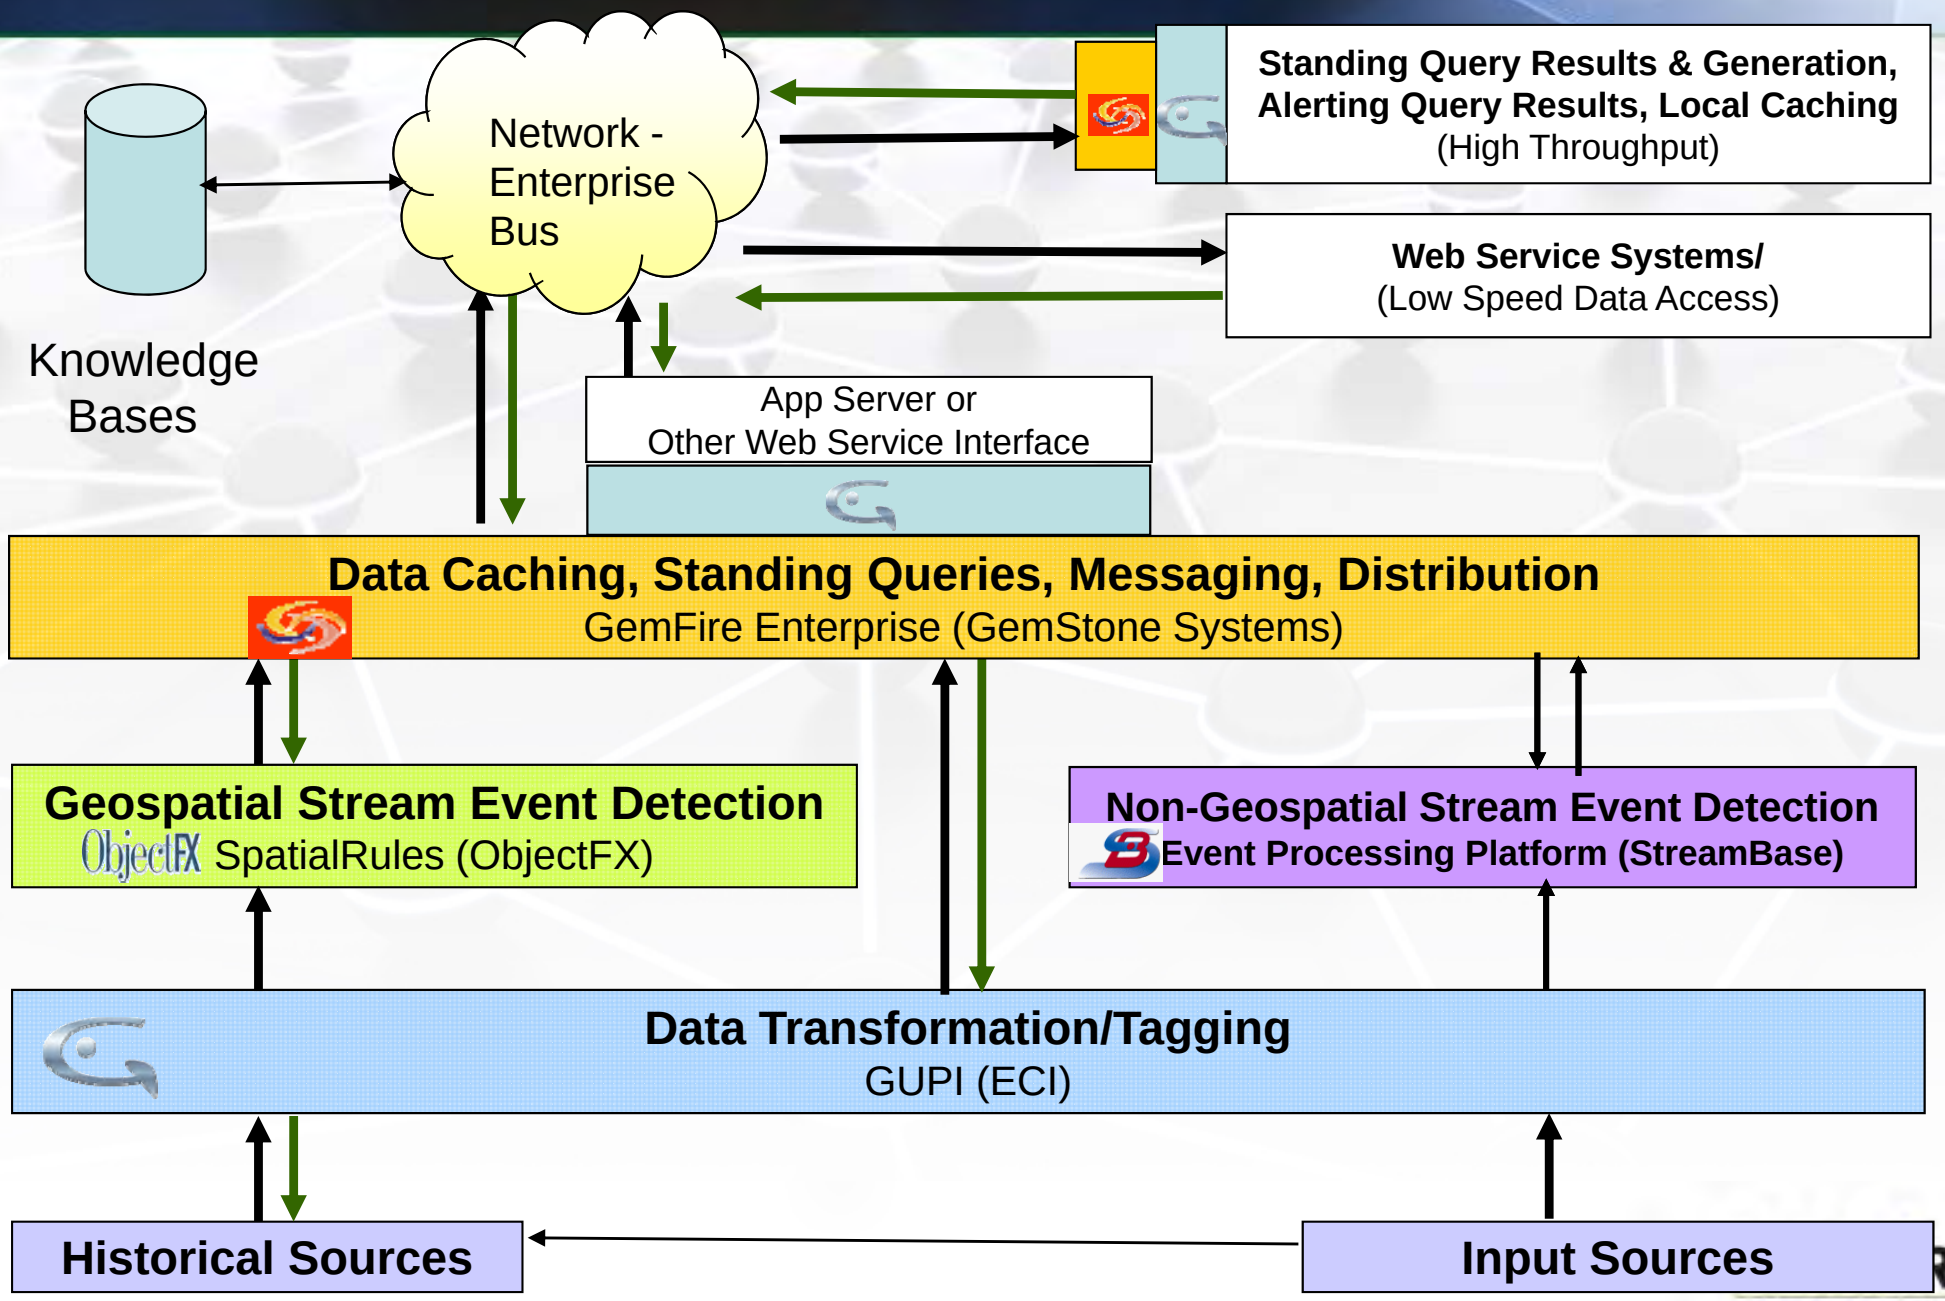

GemStone Systems
Some Strategic Customers

- 24 years in 24x7x365 global support of highly mission-critical systems.
- Over 200 Global Customers
- Over 50,000+ computers running GemFire … !

GEMFIRE
the enterprise data fabric

Why YAPOM ? Yet Another Piece of Middleware?

GEMSTONE

- **Execution Excellence**
  - Make your applications run 4 to 40 times faster
  - Ingest / Digest / Distribute vast amounts of data with extremely low latency
- **Higher ROI of IT investments & Lower Cost Per Transaction**
  - A 4x Performance Increase is like getting **3 free computers** that take **no space or electricity**
- **Enables applications to survive network outages and distressed networks**
- **Data Awareness – especially across applications/systems**
- **Supports High Availability, Fault Tolerance, and Site Failover with zero additional design/development costs**